

Junos[®] OS

MPLS Applications User Guide

Published
2019-12-26

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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

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

Junos[®] OS MPLS Applications User Guide

Copyright © 2019 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 <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation | lxxv

Documentation and Release Notes | lxxv

Using the Examples in This Manual | lxxv

 Merging a Full Example | lxxvi

 Merging a Snippet | lxxvii

Documentation Conventions | lxxvii

Documentation Feedback | lxxx

Requesting Technical Support | lxxx

 Self-Help Online Tools and Resources | lxxxi

 Creating a Service Request with JTAC | lxxxi

1

MPLS Overview

Introduction to MPLS | 3

MPLS Overview | 3

 Why Use MPLS? | 4

 Why Not Use MPLS? | 5

 How Do I Configure MPLS? | 5

 Configure the MPLS LER (Ingress) Switch and the Egress Switch | 5

 Configure LSRs for MPLS | 6

 What Does the MPLS Protocol Do? | 6

 How Does MPLS Interface to Other Protocols? | 7

 If I Have Used Cisco MPLS, What Do I Need to Know? | 7

MPLS Overview | 8

 Label Switching | 9

 Label-Switched Paths | 9

 Label-Switching Routers | 10

 Labels | 11

 Label Operations | 11

 Penultimate Hop Popping | 12

LSP Establishment	12
Static LSPs	12
Dynamic LSPs	12
MPLS Overview for ACX Series Universal Metro Routers	13
MPLS for EX Series Switches Overview	14
Benefits of MPLS	15
Additional Benefits of MPLS and Traffic Engineering	15
MPLS Applications	16
Link-Layer Support in MPLS	16
Supported MPLS Scaling Values	16
Supported MPLS Standards	20
IP and MPLS Packets on Aggregated Interfaces	23
BGP Destinations	23
IGP and BGP Destinations	25
Install Prefix IGP Overview	26
MPLS Feature Support on QFX Series and EX4600 Switches	26
Supported Features	27
Understanding MPLS Components for QFX Series and EX4600 Switches	38
Provider Edge Switches	38
MPLS Protocol and Label-Switched Paths	39
IP Over MPLS for Customer Edge Interfaces	39
BGP Layer 3 VPN Configuration	39
Routing Instances for Layer 3 VPN	39
Routing Instances for Layer 2 VPN and Layer 3 VPN	39
Ethernet Encapsulation for Layer 2 VPN	40
Provider Switch	40
Components Required for All Switches in the MPLS Network	40
Interior Gateway Protocol	40
Traffic Engineering	41
MPLS Protocol	41
RSVP	41
Family mpls	41

Understanding MPLS and Path Protection on EX Series Switches | 43

MPLS Limitations on QFX Series and EX4600 Switches | 43

MPLS Limitations on QFX10000 Switches | 44

MPLS Limitations on EX4600, EX4650, QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 Switches | 45

MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric Switches | 47

MPLS Limitations on QFX3500 Switches | 48

MPLS Configuration Overview | 48

MPLS Configuration Guidelines | 49

TTL Processing on Incoming MPLS Packets | 51

2

Configuring MPLS and Associated Features

Configuring MPLS | 55

Configuring MPLS | 55

Example: Enabling MPLS | 56

Example: Configuring MPLS on EX8200 and EX4500 Switches | 59

Verifying That MPLS Is Working Correctly | 79

Verifying the Physical Layer on the Switches | 79

Verifying the Routing Protocol | 80

Verifying the Core Interfaces Being Used for the MPLS Traffic | 80

Verifying RSVP | 81

Configuring MPLS on Provider and Provider Edge Devices | 83

Configuring MPLS on Provider Edge Switches | 83

Configuring the Ingress PE Switch | 84

Configuring the Egress PE Switch | 86

Configuring MPLS on Provider Switches | 88

Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) | 89

Configuring the Ingress PE Switch | 90

Configuring the Egress PE Switch | 92

Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect (CLI Procedure) | 95

Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure) | 99

Configuring Bidirectional Forwarding Detection (BFD) for MPLS | 101

Configuring Bidirectional Forwarding Detection for MPLS (CLI Procedure) | 101

Configuring BFD on Provider Edge and Provider Switches for an LDP-Based LSP | 102

Configuring BFD on Provider Edge and Provider Switches for an RSVP-Based LSP | 105

BFD-Triggered Local Repair for Rapid Convergence | 106

Understanding BFD-Triggered Local Protection | 107

Purpose of BFD-Triggered Local Repair | 107

Configuring BFD-Triggered Local Repair | 108

Disabling BFD-Triggered Local Repair | 108

Configuring BFD for MPLS IPv4 LSPs | 109

Configuring BFD for RSVP-Signaled LSPs | 110

Configuring a Failure Action for the BFD Session on an RSVP LSP | 112

Configuring Firewall Filters, System Log Messages, and SNMP for MPLS | 115

Configuring MPLS Firewall Filters and Policers on Switches | 115

Configuring an MPLS Firewall Filter | 116

Applying an MPLS Firewall Filter to an MPLS Interface | 116

Applying an MPLS Firewall Filter to a Loopback Interface | 117

Configuring Policers for LSPs | 118

Configuring MPLS Firewall Filters and Policers on Routers | 119

Configuring MPLS Firewall Filters | 119

Examples: Configuring MPLS Firewall Filters | 120

Configuring Policers for LSPs | 121

LSP Policer Limitations | 122

Example: Configuring an LSP Policer | 123

Configuring Automatic Policers | 124

Configuring Automatic Policers for LSPs | 125

Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs | 125

Configuring Automatic Policers for Point-to-Multipoint LSPs | 126

Disabling Automatic Policing on an LSP | 126

Example: Configuring Automatic Policing for an LSP | 127

Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets | 127

System Log Messages and SNMP Traps for MPLS | 128

Configuring Graceful Restart for MPLS | 131

Configuring MPLS-Signaled LSPs to Use GRE Tunnels | 131

Example: Configuring MPLS-Signaled LSPs to Use GRE Tunnels | 131

Graceful Restart and MPLS-Related Protocols | 133

LDP | 133

RSVP | 134

CCC and TCC | 134

Configuring IP-Based Filtering for MPLS | 135

Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic | 135

IP-Based Filtering of MPLS Traffic | 135

Selective Port Mirroring of MPLS Traffic | 137

Sample Configurations | 138

IP-Based Filtering Configuration | 138

Selective Port Mirroring Configuration | 140

Mirrored Destination Configuration | 141

Configuring Link, Node, and Path Protection for MPLS | 143

Node-Link Protection Overview | 143

Path Protection Overview | 145

Configuring Path Protection in an MPLS Network (CLI Procedure) | 145

Configuring the Primary Path | 147

Configuring the Secondary Path | 149

Configuring the Revert Timer | 149

Preventing Use of a Path That Previously Failed	150
Configuring MPLS Inter-AS Link-Node Protection with Labeled BGP	151
Understanding MPLS Inter-AS Link Protection	151
Example: Configuring MPLS Inter-AS Link-Node Protection	153
Configuring Egress Protection Service Mirroring for BGP Signaled Layer 2 Services	171
Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services	176
Example: Configuring Layer 3 VPN Egress Protection with PLR as Protector	198
Verifying Path Protection in an MPLS Network	234
Verifying the Primary Path	234
Verifying the RSVP-Enabled Interfaces	235
Verifying a Secondary Path	236
Configuring MPLS Load Balancing and Statistics 	239
MPLS Encapsulated Payload Load-balancing Overview	239
Configuring MPLS Encapsulated Payload for Load Balancing	241
Configuring MPLS to Gather Statistics	242
On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview	243
Importance of Measuring Packet Loss and Delay	243
Defining Packet Loss, Delay, and Throughput	244
Packet Loss and Delay Measurement Mechanisms	244
Packet Loss and Delay Metrics	245
Packet Loss and Delay Measurement Concepts	245
Packet Loss and Delay Measurement Functionality	248
Packet Loss and Delay Features	249
Example: Configuring On-Demand Loss and Delay Measurement	250
Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs	262
Configuring On-Demand Loss and Delay Measurement	272
Configuring Pro-Active Loss and Delay Measurements	273
Configuring Shared Risk Link Group (SRLG) 	277
SRLG Overview	277
Example: Configuring SRLG	278
Example: Excluding SRLG Links Completely for the Secondary LSP	290
Example: Configuring SRLG with Link Protection	298

Example: Configuring SRLG with Link Protection with the exclude-srlg Option | 325

Configuring MPLS Tunnels | 353

Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks | 353

Configuring IPv6 Tunneling for MPLS | 365

Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels | 367

Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview | 385

Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels | 388

Next-Hop-Based Dynamic Tunnel Localization Overview | 401

Benefits of Next-Hop-Based Dynamic Tunnel Localization | 402

Use Cases for Next-Hop-Based Dynamic Tunnel Localization | 402

Traffic Handling with Localization of Next-Hop-Based Dynamic Tunnels | 402

Configuring Next-Hop-Based Dynamic Tunnels Localization | 403

Configuring Localization for New Next-Hop-Based Dynamic Tunnels | 403

Configuring Localization for Existing Next-Hop-Based Dynamic Tunnels | 405

Troubleshooting Localized Next-Hop-Based Dynamic Tunnels | 406

Unsupported Features for Next-Hop-Based Dynamic Tunnels Localization | 407

Configuring Policy-Based Multipath Routes | 409

Policy-Based Multipath Routes Overview | 409

Understanding Policy-Based Multipath Routes | 409

Benefits of Policy-Based Multipath Routes | 410

Policy-Based Multipath Routes for Route Resolution | 410

Sample Route Resolution Using Policy-Based Multipath Routes | 411

Enhancement to Class-of-Service (CoS) Forwarding-Policy | 413

Enhancements to Policy Match Protocol | 414

Impact of Configuring Policy-Based Multipath Route on Network Performance | 415

MPLS Label-Switched Paths

MPLS Label Operations | 419

MPLS Label Overview | 419

MPLS Label Allocation | 420

Operations on MPLS Labels | 421

Understanding MPLS Label Operations | 422

- MPLS Label-Switched Paths and MPLS Labels | 422

- Reserved Labels | 423

- MPLS Label Operations | 423

- Penultimate-Hop Popping and Ultimate-Hop Popping | 425

Understanding MPLS Label Manager | 425

Understanding MPLS Label Operations on EX Series Switches | 426

- MPLS Label-Switched Paths and MPLS Labels on the Switches | 427

- Reserved Labels | 428

- MPLS Label Operations on the Switches | 428

- Penultimate-Hop Popping and Ultimate-Hop Popping | 429

How a Packet Travels Along an LSP | 430

Types of LSPs | 431

Scope of LSPs | 431

Special MPLS Labels | 432

Entropy Label Support in Mixed Mode Overview | 433

Abstract Hops for MPLS LSPs Overview | 433

- Understanding Abstract Hops | 434

- Benefits of Using Abstract Hops | 435

- Specifying a Sequence of Constraint Combinations | 435

- Avoiding New Network Configuration on Transit Nodes | 435

- Combining Centralized and Distributed Path Computation Paradigms | 436

Junos OS Implementation of Abstract Hops | 437

- Defining Abstract Hops | 437

- Using Abstract Hops in Path Constraint | 441

- Path Computation and Backtracking | 445

Sample Backtracking | 445

Example: Configuring Abstract Hops for MPLS LSPs | 446

Configuring the Maximum Number of MPLS Labels | 467

Configuring MPLS to Pop the Label on the Ultimate-Hop Router | 470

Advertising Explicit Null Labels to BGP Peers | 471

MPLS LSP Routes | 473

MPLS and Routing Tables | 474

MPLS and Traffic Protection | 476

Fast Reroute Overview | 477

Configuring Fast Reroute | 479

Detour Merging Process | 481

Detour Computations | 482

Fast Reroute Path Optimization | 482

Configuring the Optimization Interval for Fast Reroute Paths | 483

Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table | 483

Constrained-Path LSP Computation | 485

How CSPF Selects a Path | 486

CSPF Path Selection Tie-Breaking | 487

Computing CSPF Paths Offline | 488

Configuring CSPF Tie Breaking | 489

Disabling Constrained-Path LSP Computation | 489

Configuring Load Balancing Based on MPLS Labels | 490

Configuring Load Balancing Based on MPLS Labels on ACX Series Routers | 496

Example: Configuring a Constrained-Path LSP for Which Junos OS Makes All Forwarding Decisions | 500

Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and Considers Hop Constraints | 501

Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and the Secondary Path Is Explicit | 501

Path Computation for LSPs on an Overloaded Router | 503

Computing Backup Paths for LSPs Using Fate Sharing | 504

Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts | 504

Enabling IGP Shortcuts | 506

LSPs Qualified in IGP Shortcut Computations | 506

IGP Shortcut Applications | 506

IGP Shortcuts and Routing Tables | 507

IGP Shortcuts and VPNs | 508

Advertising LSPs into IGP | 508

Selecting a Forwarding LSP Next Hop | 509

Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes | 510

ECMP Flow-Based Forwarding on ACX Series Routers | 511

MPLS LSP Routers | 513

Routers in an LSP | 513

Configuring the Ingress and Egress Router Addresses for LSPs | 514

- Configuring the Ingress Router Address for LSPs | 514
- Configuring the Egress Router Address for LSPs | 514
- Preventing the Addition of Egress Router Addresses to Routing Tables | 515

Configuring the Ingress Router for MPLS-Signaled LSPs | 516

- Creating Named Paths | 516
 - Examples: Creating Named Paths | 518
- Configuring Alternate Backup Paths Using Fate Sharing | 518
 - Configuring Fate Sharing | 519
 - Implications for CSPF | 520
 - Implications for CSPF When Fate Sharing with Bypass LSPs | 520
 - Example: Configuring Fate Sharing | 521

Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs | 521

Configuring the Connection Between Ingress and Egress Routers | 522

Pinging LSPs | 522

- Pinging MPLS LSPs | 523
- Pinging Point-to-Multipoint LSPs | 523
- Pinging the Endpoint Address of MPLS LSPs | 524
- Pinging CCC LSPs | 524
- Pinging Layer 3 VPNs | 524
- Support for LSP Ping and Traceroute Commands Based on RFC 4379 | 524

Configuring MPLS LSPs | 525

Configuring LSP Metrics | 526

- Configuring Dynamic LSP Metrics | 526**

- Configuring Static LSP Metrics | 527**

Configuring a Text Description for LSPs | 528

Configuring MPLS Soft Preemption | 530

Configuring Priority and Preemption for LSPs | 531

Configuring Administrative Groups for LSPs | 532

Configuring Extended Administrative Groups for LSPs | 534

Configuring Preference Values for LSPs | 537

Disabling Path Route Recording by LSPs | 537

Achieving a Make-Before-Break, Hitless Switchover for LSPs | 538

- Specifying the Amount of Time the Router Waits to Switch Over to New Paths | 539**

- Specifying the Amount of Time to Delay the Tear Down of Old Paths | 539**

- Achieving a Hitless, MBB Switchover Without Artificial Delays | 540**

Optimizing Signaled LSPs | 541

Configuring the Smart Optimize Timer for LSPs | 545

Limiting the Number of Hops in LSPs | 546

Configuring the Bandwidth Value for LSPs | 546

Automatic Bandwidth Allocation for LSPs | 547

Configuring Automatic Bandwidth Allocation for LSPs | 547

- Configuring Automatic Bandwidth Allocation on LSPs | 549**

- Configuring the Automatic Bandwidth Allocation Interval | 550**

- Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth | 550**

- Configuring the Automatic Bandwidth Adjustment Threshold | 551**

- Configuring a Limit on Bandwidth Overflow and Underflow Samples | 552**

- Configuring Passive Bandwidth Utilization Monitoring | 554**

- Requesting Automatic Bandwidth Allocation Adjustment | 555**

Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs | 556

Configuring an LSP Across ASs | 561

Disabling Normal TTL Decrementing | 562

Adaptive LSP Configuration | 564

Damping Advertisement of LSP State Changes | 565

Configuring Primary and Secondary LSPs	565
Configuring Primary and Secondary Paths for an LSP	566
Configuring the Revert Timer for LSPs	567
Specifying the Conditions for Path Selection	567
Configuring Hot Standby of Secondary Paths for LSPs	568
Configuring Corouted Bidirectional LSPs	570
Configuring the Entropy Label for LSPs	573
Example: Configuring an Entropy Label for a BGP Labeled Unicast LSP	575
Configuring Ultimate-Hop Popping for LSPs	599
Configuring Static LSPs	603
Configuring the Ingress Router for Static LSPs	604
Example: Configuring the Ingress Router	605
Configuring the Intermediate (Transit) and Egress Routers for Static LSPs	607
Example: Configuring an Intermediate Router	608
Example: Configuring an Egress Router	609
Configuring a Bypass LSP for the Static LSP	610
Configuring the Protection Revert Timer for Static LSPs	610
Configuring Static Unicast Routes for Point-to-Multipoint LSPs	611
Configuring Static Label Switched Paths for MPLS (CLI Procedure)	612
Configuring the Ingress PE Switch	613
Configuring the Provider and the Egress PE Switch	614
Enabling Distributed CSPF for Segment Routing LSPs	615
Distributed CSPF Computation Constraints	616
Distributed CSPF Computation Algorithm	616
Label Stack Compression Enabled	617
Label Stack Compression Disabled	617
Distributed CSPF Computation Database	617
Configuring Distributed CSPF Computation Constraints	617
Distributed CSPF Computation	619
Interaction Between Distributed CSPF Computation and SRTE Features	619
Weights Associated With Paths of an SRTE Policy	619
BFD Liveliness Detection	619
inherit-label-nexthops	620
Auto-Translate Feature	620

Distributed CSPF Computation Sample Configurations | 620

Example 1 | 620

Example 2 | 621

Example 3 | 621

Color-Based Mapping of VPN Services Overview | 622

VPN Service Coloring | 623

Egress Color Assignment | 623

Ingress Color Assignment | 625

Specifying VPN Service Mapping Mode | 626

Color-IP Protocol Next Hop Resolution | 627

Fallback to IP Protocol Next Hop Resolution | 628

Supported and Unsupported Features for Color-Based Mapping of VPN Services | 628

Configuring Static Label Switched Paths for MPLS | 629

Configuring the Ingress PE Switch | 630

Configuring the Provider and the Egress PE Switch | 631

Static Segment Routing Label Switched Path | 632

Understanding Static Segment Routing LSP in MPLS Networks | 632

Introduction to Segment Routing LSPs | 633

Benefits of using Segment Routing LSPs | 634

Colored Static Segment Routing LSP | 634

Non-Colored Static Segment Routing LSP | 635

Static Segment Routing LSP Provisioning | 642

Static Segment Routing LSP Limitations | 642

Dynamic Creation of Segment Routing LSPs | 643

Color-Based Mapping of VPN Services | 649

Tunnel Templates for PCE-Initiated Segment Routing LSPs | 656

Example: Configuring Static Segment Routing Label Switched Path | 657

Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution | 676

Understanding RE-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution | 677

S-BFD Static Segment-Routing Tunnels for First-Hop Labels | 677

Limitations | 679

Configuring RE-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution | 679

Example | 681

Verify That LSPs Are Configured for Static Segment-Routing Tunnels and That S-BFD Session Status Is Visible | 683

Verify the Segment-Routing Tunnel Route with a Primary Next Hop and a Secondary Next Hop | 684

Verify the S-BFD Session of the Primary Path | 685

Configuring Explicit-Path LSPs | 686

Example: Configuring an Explicit-Path LSP | 687

Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP | 688

Configuring Point-to-Multipoint LSPs | 691

Point-to-Multipoint LSPs Overview | 691

Understanding Point-to-Multipoint LSPs | 693

Point-to-Multipoint LSP Configuration Overview | 694

Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP | 695

Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs | 720

Configuring the Primary Point-to-Multipoint LSP | 721

Configuring a Branch LSP for Point-to-Multipoint LSPs | 721

Configuring the Branch LSP as a Dynamic Path | 722

Configuring the Branch LSP as a Static Path | 722

Configuring Inter-Domain Point-to-Multipoint LSPs | 722

Configuring Link Protection for Point-to-Multipoint LSPs | 724

Configuring Graceful Restart for Point-to-Multipoint LSPs | 724

Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs | 725

 Example: Configuring Multicast RPF Check Policy for a Point-to-Multipoint LSP | 726

Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs | 726

Enabling Point-to-Point LSPs to Monitor Egress PE Routers | 727

Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases | 728

Configuring Container LSPs | 729

Dynamic Bandwidth Management Using Container LSP Overview | 729

 Understanding RSVP Multipath Extensions | 730

 Junos OS RSVP Multipath Implementation | 731

 Current Traffic Engineering Challenges | 731

 Using Container LSP as a Solution | 734

 Accommodating the New Demand X | 735

 Creating New LSPs to Meet Demand X | 735

 Assigning Bandwidth to the New LSPs | 735

 Controlling the LSP Paths | 735

 Junos OS Container LSP Implementation | 736

 Container LSP Terminology | 736

 LSP Splitting | 737

 LSP Merging | 740

 Node and Link Protection | 741

 Naming Convention | 742

 Normalization | 742

 Constraint-Based Routing Path Computation | 748

 Sampling | 749

 Support for NSR, IPG-FA, and Static Routes | 749

 Configuration Statements Supported for Container LSPs | 753

 Impact of Configuring Container LSPs on Network Performance | 757

 Supported and Unsupported Features | 758

Example: Configuring Dynamic Bandwidth Management Using Container LSP | 759

Configuring Dynamic Bandwidth Management Using Container LSP | 791

Configuring Pop-and-Forward LSPs | 797

Pop-and-Forward LSP Configuration | 797

Benefits of RSVP-TE Pop-and-Forward LSP Tunnels | 798

Pop-and-Forward LSP Tunnel Terminology | 798

Pop-and-Forward LSP Tunnel Label and Signaling | 799

Pop-and-Forward LSP Tunnel Label Stacking | 800

Construction of Label Stack at the Ingress | 800

Auto-Delegation of Label Stack | 800

Pop-and-Forward LSP Tunnel Link Protection | 801

RSVP-TE Pop-and-Forward LSP Tunnel Supported and Unsupported Features | 802

MPLS Signalling Protocols

Configuring RSVP | 807

RSVP Overview | 809

RSVP Operation Overview | 809

Understanding the RSVP Signaling Protocol | 810

RSVP Fundamentals | 811

Bandwidth Reservation Requirement | 811

Explicit Route Objects | 811

Constrained Shortest Path First | 812

Link Coloring | 813

RSVP-TE protocol extensions for FRR | 813

Junos OS RSVP Protocol Implementation | 816

RSVP Authentication | 816

Supported RSVP Standards | 816

RSVP and IGP Hello Packets and Timers | 818

RSVP Message Types | 819

Understanding RSVP Automatic Mesh | 819

Path Messages | 820

Resv Messages | 821

PathTear Messages | 821

ResvTear Messages | 821

PathErr Messages | 822

ResvErr Messages | 822

ResvConfirm Messages	822
RSVP Reservation Styles	822
RSVP Refresh Reduction	823
MTU Signaling in RSVP	824
How the Correct MTU Is Signaled in RSVP	825
Determining an Outgoing MTU Value	826
MTU Signaling in RSVP Limitations	826
Minimum RSVP Configuration	827
Configuring RSVP and MPLS	827
Example: Configuring RSVP and MPLS	828
Configuring RSVP Interfaces	829
Configuring RSVP Refresh Reduction	829
Determining the Refresh Reduction Capability of RSVP Neighbors	831
Configuring the RSVP Hello Interval	831
Configuring RSVP Authentication	832
Configuring the Bandwidth Subscription for Class Types	833
Configuring the RSVP Update Threshold on an Interface	833
Configuring RSVP for Unnumbered Interfaces	834
Configuring RSVP Node-ID Hellos	835
Example: Configuring RSVP-Signaled LSPs	836
Example: Configuring RSVP Automatic Mesh	841
Configuring Hello Acknowledgments for Nonsession RSVP Neighbors	846
Switching LSPs Away from a Network Node	846
Configuring RSVP Setup Protection	848
Configuring Load Balancing Across RSVP LSPs	848
Configuring RSVP Automatic Mesh	850
Configuring Timers for RSVP Refresh Messages	851
Preempting RSVP Sessions	852
Configuring MTU Signaling in RSVP	853
Enabling MTU Signaling in RSVP	853
Enabling Packet Fragmentation	854
Configuring Ultimate-Hop Popping for LSPs	854
Configuring RSVP to Pop the Label on the Ultimate-Hop Router	858
Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs	859

Tracing RSVP Protocol Traffic | **860**

 Examples: Tracing RSVP Protocol Traffic | **861**

Link Protection | **864**

Multiple Bypass LSPs for Link Protection | **865**

Node Protection | **865**

Fast Reroute, Node Protection, and Link Protection | **867**

 LSP Protection Overview | **867**

 LSP Protection Types Comparison | **868**

 One-to-One Backup Implementation | **868**

 Facility Backup Implementation | **869**

Configuring Link Protection on Interfaces Used by LSPs | **871**

 Configuring Bypass LSPs | **872**

 Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass LSPs | **873**

 Configuring Administrative Groups for Bypass LSPs | **873**

 Configuring the Bandwidth for Bypass LSPs | **874**

 Configuring Class of Service for Bypass LSPs | **875**

 Configuring the Hop Limit for Bypass LSPs | **875**

 Configuring the Maximum Number of Bypass LSPs | **876**

 Disabling CSPF for Bypass LSPs | **876**

 Disabling Node Protection for Bypass LSPs | **877**

 Configuring the Optimization Interval for Bypass LSPs | **877**

 Configuring an Explicit Path for Bypass LSPs | **878**

 Configuring the Amount of Bandwidth Subscribed for Bypass LSPs | **878**

 Configuring Priority and Preemption for Bypass LSPs | **879**

Configuring Node Protection or Link Protection for LSPs | **879**

Configuring Inter-AS Node and Link Protection | **880**

RSVP Graceful Restart | **881**

RSVP Graceful Restart Standard | **881**

RSVP Graceful Restart Terminology | **881**

RSVP Graceful Restart Operation | **882**

Processing the Restart Cap Object | **883**

Configuring RSVP Graceful Restart | **883**

 Enabling Graceful Restart for All Routing Protocols | **884**

 Disabling Graceful Restart for RSVP | **884**

- Disabling RSVP Helper Mode | 885
- Configuring the Maximum Helper Recovery Time | 885
- Configuring the Maximum Helper Restart Time | 885

RSVP LSP Tunnels Overview | 885

Example: RSVP LSP Tunnel Configuration | 888

Verifying Your Work | 904

Router 0 | 905

Router 1 | 911

Configuring Link Management Protocol Peers | 912

Configuring Link Management Protocol Traffic Engineering Links | 913

Configuring Peer Interfaces in OSPF and RSVP | 914

Defining Label-Switched Paths for the FA-LSP | 915

Establishing FA-LSP Path Information | 915

Option: Tearing Down RSVP LSPs Gracefully | 916

Configuring LDP | 917

LDP Introduction | 918

Understanding the LDP Signaling Protocol | 919

Example: Configuring LDP-Signaled LSPs | 920

Supported LDP Standards | 923

Junos OS LDP Protocol Implementation | 924

LDP Operation | 924

LDP Message Types | 925

Discovery Messages | 925

Session Messages | 926

Advertisement Messages | 926

Notification Messages | 926

Discovery Messages | 927

Session Messages | 927

Advertisement Messages | 927

Notification Messages | 927

Tunneling LDP LSPs in RSVP LSPs | 928

Tunneling LDP LSPs in RSVP LSPs Overview | 928

Benefits of Tunneling LDP LSPs in RSVP LSPs | 929

Label Operations	930
LDP Session Protection	931
LDP Native IPv6 Support Overview	932
Longest Match Support for LDP Overview	933
Minimum LDP Configuration	933
Enabling and Disabling LDP	934
Configuring the LDP Timer for Hello Messages	934
Configuring the LDP Timer for Link Hello Messages	935
Configuring the LDP Timer for Targeted Hello Messages	935
Configuring the Delay Before LDP Neighbors Are Considered Down	935
Configuring the LDP Hold Time for Link Hello Messages	936
Configuring the LDP Hold Time for Targeted Hello Messages	937
Enabling Strict Targeted Hello Messages for LDP	937
Configuring the Interval for LDP Keepalive Messages	938
Configuring the LDP Keepalive Timeout	938
Configuring Longest Match for LDP	939
Example: Configuring Longest Match for LDP	940
Configuring LDP Route Preferences	959
LDP Graceful Restart	959
Configuring LDP Graceful Restart	960
Enabling Graceful Restart	960
Disabling LDP Graceful Restart or Helper Mode	961
Configuring Reconnect Time	962
Configuring Recovery Time and Maximum Recovery Time	962
Filtering Inbound LDP Label Bindings	963
Examples: Filtering Inbound LDP Label Bindings	964
Filtering Outbound LDP Label Bindings	965
Examples: Filtering Outbound LDP Label Bindings	967
Specifying the Transport Address Used by LDP	968
Control Transport Address Used for Targeted-LDP Session	969
Benefits of Controlling Transport Address Used for Targeted-LDP Session	969
Targeted-LDP Transport Address Overview	969
Transport Address Preference	970
Troubleshooting Transport Address Configuration	970

Configuring the Prefixes Advertised into LDP from the Routing Table | 971

Example: Configuring the Prefixes Advertised into LDP | 972

Configuring FEC Deaggregation | 973

Configuring Policers for LDP FECs | 974

Configuring LDP IPv4 FEC Filtering | 975

Configuring BFD for LDP LSPs | 975

Configuring ECMP-Aware BFD for LDP LSPs | 979

Configuring a Failure Action for the BFD Session on an LDP LSP | 980

Configuring the Holddown Interval for the BFD Session | 980

Configuring LDP Link Protection | 981

Example: Configuring LDP Link Protection | 983

LDP Link Protection Overview | 983

Introduction to LDP | 983

Junos OS LDP Protocol Implementation | 984

Understanding Multipoint Extensions to LDP | 984

Using Multipoint Extensions to LDP on Targeted LDP Sessions | 984

Current Limitations of LDP Link Protection | 985

Using RSVP LSP as a Solution | 986

Understanding Multicast LDP Link Protection | 988

Different Modes for Providing LDP Link Protection | 989

Label Operation for LDP Link Protection | 991

Sample Multicast LDP Link Protection Configuration | 998

Make-Before-Break | 1000

Caveats and Limitations | 1001

Example: Configuring LDP Link Protection | 1002

Understanding Multicast-Only Fast Reroute | 1014

PIM Functionality | 1016

Multipoint LDP Functionality | 1017

Packet Forwarding | 1018

Limitations and Caveats | 1019

Configuring Multicast-Only Fast Reroute | 1021

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024

Example: Configuring LDP Downstream on Demand | 1046

Configuring LDP Native IPv6 Support | 1053

Example: Configuring LDP Native IPv6 Support | **1054**

Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs | **1072**

Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs | **1072**

How M-LDP Works | **1073**

Terminology | **1078**

Ingress Join Translation and Pseudo Interface Handling | **1079**

Ingress Splicing | **1079**

Reverse Path Forwarding | **1079**

LSP Root Detection | **1080**

Egress Join Translation and Pseudo Interface Handling | **1080**

Egress Splicing | **1080**

Supported Functionality | **1080**

Unsupported Functionality | **1081**

LDP Functionality | **1081**

Egress LER Functionality | **1081**

Transit LSR Functionality | **1082**

Ingress LER Functionality | **1082**

Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs | **1082**

LDP Mapping Server for Interoperability of Segment Routing with LDP Overview | **1110**

Interoperability of Segment Routing with LDP Using OSPF | **1110**

Interoperability of Segment Routing with LDP Using ISIS | **1112**

Configuring Miscellaneous LDP Properties | **1115**

Configuring LDP to Use the IGP Route Metric | **1116**

Preventing Addition of Ingress Routes to the inet.0 Routing Table | **1116**

Multiple-Instance LDP and Carrier-of-Carriers VPNs | **1116**

Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router | **1117**

Enabling LDP over RSVP-Established LSPs | **1117**

Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks | **1118**

Configuring the TCP MD5 Signature for LDP Sessions | **1118**

Configuring LDP Session Protection | **1120**

Disabling SNMP Traps for LDP | **1120**

Configuring LDP Synchronization with the IGP on LDP Links | **1121**

Configuring LDP Synchronization with the IGP on the Router | **1121**

Configuring the Label Withdrawal Timer | **1121**

Ignoring the LDP Subnet Check | 1122

Configuring LDP LSP Traceroute | 1123

Collecting LDP Statistics | 1124

LDP Statistics Output | 1125

Disabling LDP Statistics on the Penultimate-Hop Router | 1126

LDP Statistics Limitations | 1126

Tracing LDP Protocol Traffic | 1127

Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels | 1127

Tracing LDP Protocol Traffic Within FECs | 1128

Examples: Tracing LDP Protocol Traffic | 1129

MPLS Traffic Engineering

Understanding MPLS Traffic Engineering | 1133

MPLS and Traffic Engineering | 1133

MPLS Traffic Engineering and Signaling Protocols Overview | 1134

Traffic Engineering Capabilities | 1136

Components of Traffic Engineering | 1136

Configuring Traffic Engineering for LSPs | 1137

Using LSPs for Both BGP and IGP Traffic Forwarding | 1137

Using LSPs for Forwarding in Virtual Private Networks | 1138

Using RSVP and LDP Routes for Forwarding but Not Route Selection | 1138

Advertising the LSP Metric in Summary LSAs | 1139

Enabling Interarea Traffic Engineering | 1140

Enabling Inter-AS Traffic Engineering for LSPs | 1141

Inter-AS Traffic Engineering Requirements | 1141

Inter-AS Traffic Engineering Limitations | 1142

Configuring OSPF Passive TE Mode | 1143

Packet Forwarding Component | 1143

Packet Forwarding Based on Label Swapping | 1144

How a Packet Traverses an MPLS Backbone | 1144

Information Distribution Component | 1145

Path Selection Component | 1145

Offline Path Planning and Analysis | 1146

Signaling Component | 1147

Flexible LSP Calculation and Configuration | **1147**

Link-State Distribution Using BGP Overview | **1148**

Role of an Interior Gateway Protocol | **1149**

Limitations of an Interior Gateway Protocol | **1150**

Need for Spanning Link-State Distribution | **1150**

Using BGP as a Solution | **1150**

Overview | **1151**

Implementation | **1151**

Supported and Unsupported Features | **1156**

BGP Link-State Extensions for Source Packet Routing in Networking (SPRING) | **1157**

Source Packet Routing in Networking (SPRING) | **1157**

Flow of BGP Link-State SPRING Data | **1157**

Supported BGP Link-State Attributes and TLVs, and Unsupported Features for BGP Link-State with SPRING | **1159**

Verifying NLRI Node Learned Through BGP with OSPF as IGP | **1160**

| **1160**

Verifying the Prefix NLRI Learned Through BGP with OSPF as IGP | **1161**

| **1161**

Example: Configuring Link State Distribution Using BGP | **1163**

Configuring Link State Distribution Using BGP | **1186**

Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages | **1189**

PathErr Messages | **1190**

Identifying the Problem Link | **1191**

Configuring the Router to Improve Traffic Engineering Database Accuracy | **1191**

Configuring DiffServ-Aware Traffic Engineering to Achieve Service Level Guarantees on an MPLS network | 1193

DiffServ-Aware Traffic Engineering Introduction | **1194**

DiffServ-Aware Traffic Engineering Standards | **1194**

DiffServ-Aware Traffic Engineering Terminology | **1194**

DiffServ-Aware Traffic Engineering Features | **1195**

DiffServ-Aware Traffic Engineered LSPs | **1196**

DiffServ-Aware Traffic Engineered LSPs Overview | **1196**

DiffServ-Aware Traffic Engineered LSPs Operation | **1197**

Multiclass LSP Overview | **1198**

Multiclass LSPs	1198
Establishing a Multiclass LSP on the Differentiated Services Domain	1199
Configuring Routers for DiffServ-Aware Traffic Engineering	1199
Configuring the Bandwidth Model	1201
Configuring Traffic Engineering Classes	1202
Requirements and Limitations for the Traffic Engineering Class Matrix	1203
Configuring Class of Service for DiffServ-Aware Traffic Engineering	1204
LSP Bandwidth Oversubscription Overview	1204
LSP Size Oversubscription	1205
LSP Link Size Oversubscription	1206
Class Type Oversubscription and Local Oversubscription Multipliers	1206
Class Type Bandwidth and the LOM	1207
LOM Calculation for the MAM and Extended MAM Bandwidth Models	1207
LOM Calculation for the Russian Dolls Bandwidth Model	1208
Example: LOM Calculation	1208
Configuring the Bandwidth Subscription Percentage for LSPs	1210
Constraints on Configuring Bandwidth Subscription	1211
Configuring LSPs for DiffServ-Aware Traffic Engineering	1212
Configuring Class of Service for the Interfaces	1213
Configuring IGP	1213
Configuring Traffic-Engineered LSPs	1213
Configuring Policing for LSPs	1214
Configuring Fast Reroute for Traffic-Engineered LSPs	1214
Configuring Multiclass LSPs	1215
Configuring Class of Service for the Interfaces	1216
Configuring the IGP	1216
Configuring Class-Type Bandwidth Constraints for Multiclass LSPs	1216
Configuring Policing for Multiclass LSPs	1218
Configuring Fast Reroute for Multiclass LSPs	1218

MPLS Transport Profile

Configuring Operation, Administration, and Maintenance (OAM) for MPLS | 1221

Configuring the MPLS Transport Profile for OAM | 1221

MPLS Transport Profile Overview | 1221

Example: Configuring the MPLS Transport Profile for OAM | 1222

Configuring OAM Ingress Policies for LDP | 1238

Tracing MPLS and LSP Packets and Operations | 1239

Configuring MPLS Pseudowires | 1241

Ethernet Pseudowire Overview | 1241

Example: Ethernet Pseudowire Base Configuration | 1242

Pseudowire Overview for ACX Series Universal Metro Routers | 1246

Understanding Multisegment Pseudowire for FEC 129 | 1248

Understanding Multisegment Pseudowire | 1248

Using FEC 129 for Multisegment Pseudowire | 1249

Establishing a Multisegment Pseudowire Overview | 1250

Pseudowire Status Support for Multisegment Pseudowire | 1250

Pseudowire Status Behavior on T-PE | 1250

Pseudowire Status Behavior on S-PE | 1251

Pseudowire TLV Support for MS-PW | 1251

Supported and Unsupported Features | 1252

Example: Configuring a Multisegment Pseudowire | 1253

MPLS Stitching For Virtual Machine Connection | 1299

When Would I Use Stitching? | 1300

How Does MPLS Stitching Work? | 1300

How Do I Configure Stitching? | 1301

Which Switches Support Stitching? | 1301

Q&A | 1302

TDM Pseudowires Overview | 1302

Example: TDM Pseudowire Base Configuration | 1303

Configuring Load Balancing for Ethernet Pseudowires | 1308

Configuring Load Balancing Based on MAC Addresses | 1309

Configuring Class-of-Service (CoS) for MPLS | 1311

Configuring Class of Service for MPLS LSPs | 1311

Class of Service for MPLS Overview | **1312**

Configuring the MPLS CoS Values | **1312**

Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value | **1314**

Configuring MPLS Rewrite Rules | 1315

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet | **1315**

Rewriting MPLS and IPv4 Packet Headers | **1316**

Configuring CoS Bits for an MPLS Network (CLI Procedure) | 1317

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure) | 1318

Configuring CoS | **1318**

Configuring an LSP Policer | **1319**

Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) | 1320

Configuring CoS | **1321**

Configuring an LSP Policer | **1322**

Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) | 1323

Understanding Using CoS with MPLS Networks on EX Series Switches | 1324

EXP Classifiers and EXP rewrite Rules | **1325**

Guidelines for Using CoS Classifiers on CCCs | **1325**

Using CoS Classifiers with IP over MPLS | **1326**

Setting CoS Bits in an MPLS Header | **1326**

EXP Rewrite Rules | **1328**

Policer | **1328**

Schedulers | **1329**

Example: Combining CoS with MPLS on EX Series Switches | 1329

Understanding CoS MPLS EXP Classifiers and Rewrite Rules | 1345

EXP Classifiers | **1346**

EXP Rewrite Rules | **1348**

Schedulers | **1348**

Configuring Rewrite Rules for MPLS EXP Classifiers | 1349

Configuring CoS Bits for an MPLS Network | 1351

Configuring a Global MPLS EXP Classifier | 1352

Configuring Generalized MPLS (GMPLS) | 1353

Introduction to GMPLS | 1353

GMPLS Terms and Acronyms | 1355

Supported GMPLS Standards | 1356

GMPLS Operation | 1357

GMPLS and OSPF | 1358

GMPLS and CSPF | 1358

GMPLS Features | 1359

LMP Configuration Overview | 1359

Configuring LMP Traffic Engineering Links | 1361

Configuring the Local IP Address for Traffic Engineering Links | 1361

Configuring the Remote IP Address for Traffic Engineering Links | 1362

Configuring the Remote ID for Traffic Engineering Links | 1362

Configuring LMP Peers | 1363

Configuring the ID for LMP Peers | 1364

Configuring the Interface for Control Channels Between LMP Peers | 1364

Configuring the LMP Control Channel Interface for the Peer | 1365

Configuring the Remote IP Address for LMP Control Channels | 1366

Configuring Hello Message Intervals for LMP Control Channels | 1366

Controlling Message Exchange for LMP Control Channels | 1367

Preventing the Local Peer from Initiating LMP Negotiation | 1367

Associating Traffic Engineering Links with LMP Peers | 1368

Disabling the Traffic Engineering Link for LMP Peers | 1368

Configuring RSVP and OSPF for LMP Peer Interfaces | 1369

Configuring RSVP Signaling for LMP Peer Interfaces | 1369

Configuring OSPF Routing for LMP Peer Interfaces | 1370

Configuring the Hello Interval for LMP Peer Interfaces | 1370

Configuring MPLS Paths for GMPLS | 1371

Tracing LMP Traffic | 1371

Configuring MPLS LSPs for GMPLS | 1372

Configuring the Encoding Type | 1373

Configuring the GPID | 1373

Configuring the Signal Bandwidth Type | 1374

Configuring GMPLS Bidirectional LSPs | 1374

Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS | 1374

Gracefully Tearing Down GMPLS LSPs | 1375

Temporarily Deleting GMPLS LSPs | 1375

Permanently Deleting GMPLS LSPs | 1376

Configuring the Graceful Deletion Timeout Interval | 1376

GMPLS RSVP-TE VLAN LSP Signaling Overview | 1377

Understanding GMPLS RSVP-TE Signaling | 1377

Need for GMPLS RSVP-TE VLAN LSP Signaling | 1378

GMPLS RSVP-TE VLAN LSP Signaling Functionality | 1379

LSP Hierarchy with GMPLS RSVP-TE VLAN LSP | 1380

Path Specification for GMPLS RSVP-TE VLAN LSP | 1380

GMPLS RSVP-TE VLAN LSP Configuration | 1380

Associated Bidirectional Packet LSP | 1382

Make-Before-Break for Associated Bidirectional Packet and GMPLS RSVP-TE VLAN LSP | 1382

Supported and Unsupported Features | 1383

Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling | 1384

7

MPLS BGP VPNs

Configuring MPLS VPNs | 1415

MPLS VPN Overview | 1416

MPLS VPN Topology | 1416

MPLS VPN Routing | 1418

VRF Instances | 1418

Route Distinguishers | 1418

MPLS Exceptions on SRX Series Devices | 1419

Understanding IPv6 Layer 3 VPNs | 1420

Understanding Using MPLS-Based Layer 3 VPNs on Switches | 1420

MPLS-Based Layer 3 VPNs | 1420

Configuring a BGP Session for MPLS VPNs (CLI Procedure) | 1421

Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure) | 1423

Configuring Routing Options for MPLS VPNs (CLI Procedure) | 1423

Configuring a Routing Instance for MPLS VPNs (CLI Procedure) | 1424

Chained Composite Next Hops for Transit Devices for VPNs | 1426

Understanding MPLS Layer 2 VPNs | 1427

Understanding Ethernet-over-MPLS (L2 Circuit)	1428
Ethernet-over-MPLS in Data Centers	1428
MPLS Layer 2 VPN Configuration Overview	1429
Configuring a Routing Policy for MPLS Layer 2 VPNs (CLI Procedure)	1430
Configuring Interfaces for Layer 2 VPNs (CLI Procedure)	1432
Configuring Ethernet over MPLS (L2 Circuit)	1433
Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)	1435
Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)	1436
Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit	1436
Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit	1437
Example: Configuring MPLS-Based Layer 2 VPNs	1438
Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches	1456
MPLS-Based Layer 2 VPNs	1456
Layer 2 Circuits	1457
MPLS-Based Layer 3 VPNs	1457
Comparing an MPLS-Based Layer 2 VPN and an MPLS-Based Layer 3 VPN	1458
Verifying an MPLS Layer 2 VPN Configuration	1459
Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)	1460
Understanding MPLS Layer 2 Circuits	1463
MPLS Layer 2 Circuit Configuration Overview	1464
Configuring an MPLS Layer 2 Circuit (CLI Procedure)	1465
Verifying an MPLS Layer 2 Circuit Configuration	1466
Configuring an IGP and the LDP Signaling Protocol (CLI Procedure)	1466
Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)	1468
Understanding MPLS Layer 3 VPNs	1471
MPLS Layer 3 VPN Configuration Overview	1471
Configuring a Routing Policy for MPLS Layer 3 VPNs (CLI Procedure)	1473
Verifying an MPLS Layer 3 VPN Configuration	1474
Example: Configuring MPLS-Based Layer 3 VPNs	1474
Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks	1486
Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches	1498

Configuring an MPLS-Based Layer 3 VPN (CLI Procedure) | 1512

Configuring TCC and Layer 2.5 Switching | 1515

TCC and Layer 2.5 Switching Overview | 1515

Configuring VLAN TCC Encapsulation | 1516

Configuring Translation Cross-Connect Interface Switching | 1518

Configuring CLNS VPNs | 1521

CLNS Overview | 1521

CLNS Configuration Overview | 1522

Understanding ES-IS for CLNS | 1524

Example: Configuring ES-IS for CLNS | 1525

Understanding IS-IS for CLNS | 1527

Example: Configuring IS-IS for CLNS | 1528

Understanding Static Routes for CLNS | 1531

Example: Configuring Static Routes for CLNS When No IGP is Present | 1532

Understanding BGP for CLNS VPNs | 1535

Example: Configuring BGP for CLNS VPNs | 1535

Example: Configuring a VPN Routing Instance for CLNS | 1538

Verifying a CLNS VPN Configuration | 1541

Configuring VPLS | 1545

VPLS Overview | 1546

Sample VPLS Topology | 1546

VPLS on PE Routers | 1547

Using an Ethernet Switch as the VPLS CE Device | 1549

VPLS Exceptions on SRX Series Devices | 1549

VPLS Configuration Overview | 1550

Migrating from FEC128 LDP-VPLS to EVPN Overview | 1552

Technology Overview and Benefits | 1552

FEC128 LDP-VPLS to EVPN Migration | 1553

Sample Configuration for LDP-VPLS to EVPN Migration | 1555

LDP-VPLS Configuration | 1555

EVPN Migration Configuration | 1557

Reverting to VPLS | 1559

LDP-VPLS to EVPN Migration and Other Features | 1560

Understanding VPLS Interfaces | 1561

Interface Name | 1562

Encapsulation Type | 1562

Flexible VLAN Tagging | 1562

VLAN Rewrite | 1563

Example: Configuring Routing Interfaces on the VPLS PE Router | 1564

Example: Configuring the Interface to the VPLS CE Device | 1565

VPLS Filters and Policers Overview | 1567

Example: Configuring VPLS Filters | 1567

Example: Configuring VPLS Policers | 1571

Understanding VPLS Routing Instances | 1573

BGP Signaling | 1574

VPLS Routing Table | 1575

Trace Options | 1576

Example: Configuring the VPLS Routing Instance | 1576

Example: Configuring Automatic Site Identifiers for VPLS | 1580

Example: Configuring OSPF on the VPLS PE Router | 1581

Example: Configuring RSVP on the VPLS PE Router | 1583

Example: Configuring MPLS on the VPLS PE Router | 1585

Example: Configuring LDP on the VPLS PE Router | 1587

Example: Configuring VPLS over GRE with IPsec VPNs | 1588

Example: Configuring VPLS with BGP Signaling | 1610

Example: Configuring BGP on the VPLS PE Router | 1628

Example: Configuring Routing Options on the VPLS PE Router | 1630

Understanding VPLS VLAN Encapsulation | 1631

Understanding VPLS VLAN Encapsulation on a Logical Interface | 1632

Example: Configuring VPLS VLAN Encapsulation | 1633

Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces | 1637

Example: Configuring Extended VLAN VPLS Encapsulation | 1639

Configuring Circuit Cross-Connect (CCC) and Translational Cross-Connect (TCC) | 1643

CCC Overview | 1643

Understanding Carrier-of-Carriers VPNs | 1644

- Internet Service Provider as the Customer | 1646**

- VPN Service Provider as the Customer | 1646**

Understanding Interprovider and Carrier-of-Carriers VPNs | 1647

Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit (CLI Procedure) | 1648

VLAN CCC Encapsulation on Transport Side of Pseudowire Client Logical Interfaces Overview | 1651

- Pseudowire Configuration from Access Node | 1651**

- Pseudowire Configuration from Aggregation Node | 1652**

Transmitting Nonstandard BPDUs | 1654

TCC Overview | 1654

Configuring Layer 2 Switching Cross-Connects Using CCC | 1655

- Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects | 1656**

 - Configuring ATM Encapsulation for Layer 2 Switching Cross-Connects | 1657**

 - Configuring Ethernet Encapsulation for Layer 2 Switching Cross-Connects | 1657**

 - Configuring Ethernet VLAN Encapsulation for Layer 2 Switching Cross-Connects | 1658**

 - Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching Cross-Connects | 1659**

 - Configuring Frame Relay Encapsulation for Layer 2 Switching Cross-Connects | 1660**

 - Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching Cross-Connects | 1661**

- Configuring the CCC Connection for Layer 2 Switching Cross-Connects | 1661**

- Configuring MPLS for Layer 2 Switching Cross-Connects | 1662**

- Example: Configuring a Layer 2 Switching Cross-Connect | 1662**

- Configuring Layer 2 Switching Cross-Connect on ACX5440 | 1664**

Configuring MPLS LSP Tunnel Cross-Connects Using CCC | 1665

- Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects | 1667**

- Configuring the CCC Connection for LSP Tunnel Cross-Connects | 1669**

- Example: Configuring an LSP Tunnel Cross-Connect | 1669**

Configuring TCC | 1671

- Configuring the Encapsulation for Layer 2 Switching TCCs | 1671**

 - Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching TCCs | 1672**

 - Configuring ATM Encapsulation for Layer 2 Switching TCCs | 1672**

 - Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs | 1672**

 - Configuring Ethernet Encapsulation for Layer 2 Switching TCCs | 1673**

Configuring Ethernet Extended VLAN Encapsulation for Layer 2 Switching TCCs | 1674

Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations | 1675

Configuring the Connection for Layer 2 Switching TCCs | 1675

Configuring MPLS for Layer 2 Switching TCCs | 1676

CCC and TCC Graceful Restart | 1677

Configuring CCC and TCC Graceful Restart | 1678

Configuring an MPLS-Based VLAN CCC Using the Connection Method (CLI Procedure) | 1679

Configuring CCC Switching for Point-to-Multipoint LSPs | 1681

Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers | 1682

Configuring Local Receivers on a Point-to-Multipoint CCC LSP Switch on Ingress PE Routers | 1682

Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers | 1683

Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN (CLI Procedure) | 1684

Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit (CLI Procedure) | 1689

MPLS for Software Defined Network (SDN)

Introduction to Path Computation Element Protocol (PCEP) | 1695

PCEP Overview | 1695

Configuring PCEP for MPLS RSVP-TE | 1697

Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697

Understanding MPLS RSVP-TE | 1698

Current MPLS RSVP-TE Limitations | 1699

Use of an External Path Computing Entity | 1701

Components of External Path Computing | 1701

Path Computation Element | 1701

Path Computation Client | 1703

Path Computation Element Protocol | 1703

Interaction Between a PCE and a PCC Using PCEP | 1703

LSP Behavior with External Computing | 1707

LSP Types | 1707

LSP Control Mode | 1708

Configuration Statements Supported for External Computing | 1709

PCE-Controlled LSP Protection | 1709

PCE-Controlled LSP ERO | 1710

PCE-Controlled Point-to-Multipoint RSVP-TE LSPs | 1710

PCE-Initiated Point-to-Point LSPs | 1711

PCE-Initiated Bypass LSP | 1712

Understanding PCE-Initiated Bypass LSPs | 1712

Benefits of PCE-Initiated Bypass LSP | 1713

Behavior of PCE-Initiated Bypass LSPs During PCEP Session Failure | 1713

PCE-Initiated Point-to-Multipoint LSPs | 1713

Auto-Bandwidth and PCE-Controlled LSP | 1714

TCP-MD5 Authentication for PCEP Sessions | 1714

Impact of Client-Side PCE Implementation on Network Performance | 1715

Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE | 1717

Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1734

Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1746

Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Controlled Point-to-Multipoint LSPs | 1750

Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs | 1771

Benefits of PCE-Initiated Point-to-Multipoint LSPs | 1772

Signaling of PCE-Initiated Point-to-Multipoint LSPs | 1772

Behavior of PCE-Initiated Point-to-Multipoint LSPs After PCEP Session Failure | 1772

Configuring PCE-Initiated Point-to-Multipoint LSP Capability | 1773

Supported and Unsupported Features for PCE-Initiated Point-to-Multipoint LSPs | 1773

Mapping PCE-initiated Point-To-Multipoint LSPs to MVPN | 1773

Configuring PCEP for MPLS SPRING-TE | 1777

Support of SPRING-TE for the Path Computation Element Protocol Overview | 1777

SPRING for Traffic Engineering | 1777

Junos OS Implementation of PCEP for SPRING-TE LSPs | 1778

SPRING-TE Module | 1779

Traffic Engineering Database | 1779

PCEP Interaction | 1779

Configuration of PCEP for SPRING-TE | 1780

Limitations and Unsupported Features for PCEP SPRING-TE | 1781

Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs | 1782

Static Segment Routing Label Switched Path | 1811

Understanding Static Segment Routing LSP in MPLS Networks | 1812

Introduction to Segment Routing LSPs | 1812

Benefits of using Segment Routing LSPs | 1813

Colored Static Segment Routing LSP | 1813

Non-Colored Static Segment Routing LSP | 1814

Static Segment Routing LSP Provisioning | 1821

Static Segment Routing LSP Limitations | 1821

Dynamic Creation of Segment Routing LSPs | 1822

Color-Based Mapping of VPN Services | 1828

Tunnel Templates for PCE-Initiated Segment Routing LSPs | 1835

Example: Configuring Static Segment Routing Label Switched Path | 1836

9

Troubleshooting MPLS

Troubleshooting MPLS | 1859

Verify MPLS Interfaces | 1866

Verify the MPLS Configuration | 1868

Checklist for Checking the MPLS Layer | 1871

Checking the MPLS Layer | 1872

Verify the LSP | 1874

Verify the LSP Route on the Transit Router | 1878

Verify the LSP Route on the Ingress Router | 1880

Verify MPLS Labels with the traceroute Command | 1882

Verify MPLS Labels with the ping Command | 1883

Verify the MPLS Configuration | 1885

Take Appropriate Action | 1888

Verify the LSP Again | 1889

Verify That Node-Link Protection Is Up | 1893

Verify That Link Protection Is Up | 1900

Many-to-One Link Protection (Facility Backup) Overview | 1905

Verify One-to-One Backup | 1907

Verify That the Primary Path Is Operational | 1914

Verify That the Secondary Path Is Established	1916
Verify the LSP	1919
Verify the LSP Route on the Transit Router	1923
Verify the LSP Route on the Ingress Router	1924
Verify MPLS Labels with the traceroute Command	1926
Verify MPLS Labels with the ping Command	1928
Take Appropriate Action	1929
Verify the LSP Again	1931
Checklist for Working with the Layered MPLS Troubleshooting Model	1935
Understanding the Layered MPLS Troubleshooting Model	1935
Checklist for Verifying the Physical Layer	1943
Verifying the Physical Layer	1944
Verify the LSP	1946
Verify Router Connection	1948
Verify Interfaces	1949
Take Appropriate Action	1950
Verify the LSP Again	1951
Verify the LSP	1953
Verify Router Connection	1955
Verify Interfaces	1956
Take Appropriate Action	1957
Verify the LSP Again	1958
Checklist for Checking the Data Link Layer	1960
Checking the Data Link Layer	1961
Verify the LSP	1963
Verify Interfaces	1964
Take Appropriate Action	1969
Verify the LSP Again	1970
Verify the LSP	1975
Verify Interfaces	1977
Take Appropriate Action	1981
Verify the LSP Again	1983
Checklist for Verifying the IP and IGP Layers	1988
Verifying the IP and IGP Layers	1990

Verifying the IP Layer | 1992**Verify the LSP | 1993****Verify IP Addressing | 1994****Verify Neighbors or Adjacencies at the IP Layer | 1996****Take Appropriate Action | 2000****Verify the LSP Again | 2002****Verify the LSP | 2006****Verify IP Addressing | 2007****Verify Neighbors or Adjacencies at the IP Layer | 2009****Take Appropriate Action | 2014****Verify the LSP Again | 2015****Verifying the OSPF Protocol | 2019****Verify the LSP | 2020****Verify OSPF Interfaces | 2025****Verify OSPF Neighbors | 2027****Verify the OSPF Protocol Configuration | 2028****Take Appropriate Action | 2029****Verify the LSP Again | 2031****Verify the LSP | 2035****Verify OSPF Interfaces | 2039****Verify OSPF Neighbors | 2041****Verify the LSP Again | 2042****Verify the LSP | 2046****Verify IS-IS Adjacencies and Interfaces | 2048****Verify the IS-IS Configuration | 2050****Verify the LSP Again | 2051****Checklist for Checking the RSVP Layer | 2055****Checking the RSVP Layer | 2056****Verify the LSP | 2059****Verify RSVP Sessions | 2060****Verify RSVP Neighbors | 2063****Verify RSVP Interfaces | 2064****Verify the RSVP Protocol Configuration | 2066****Take Appropriate Action | 2067**

Verify the LSP Again	2069
Verify the LSP	2073
Verify RSVP Sessions	2075
Verify RSVP Neighbors	2077
Verify RSVP Interfaces	2079
Verify the RSVP Protocol Configuration	2081
Take Appropriate Action	2082
Verify the LSP Again	2083
Checklist for Determining LSP Status	2087
Determining LSP Statistics	2088
Checklist for Verifying LSP Use	2090
Verifying LSP Use in Your Network	2090
Verifying an LSP on the Ingress Router	2091
Verifying an LSP on a Transit Router	2093
Verifying an LSP on the Ingress Router	2094
Verifying an LSP on a Transit Router	2096
Verify That Load Balancing Is Working	2097
Example: Load-Balanced MPLS Network	2101
Router Configurations for the Load-Balanced MPLS Network	2102
Traffic Flows Before Load Balancing	2117
Verify the Operation of Uneven Bandwidth Load Balancing	2119
Checklist for Collecting Crash Data	2121
Understand Crash Data Collection	2123
Collect Crash Data for a Routing Engine Kernel	2124
Check the Routing Engine Core Files	2124
List the Core Files	2125
Compress the vmcore File	2126
Log Software Version Information	2127
Open a Case with JTAC	2128
Check the Routing Engine Core Files	2128
List the Core Files	2129
Compress the vmcore File	2130
Log Software Version Information	2131
Open a Case with JTAC	2132

Collect Crash Data for Routing Engine Daemons	2132
Check for Daemon Core Files	2132
List the Daemon Core Files	2134
Compress the Daemon Core Files	2135
Log Software Version Information	2136
Open a Case with JTAC	2137
Collect and Send Routing Engine Crash Data to JTAC	2137
Check for Daemon Core Files	2138
List the Daemon Core Files	2139
Compress the Daemon Core Files	2141
Collect Crash Data for the Packet Forwarding Engine Microkernel	2141
Display the Crash Stack Traceback and Registration Information	2143
Clear the NVRAM Contents	2147
Check Packet Forwarding Engine Microkernel Core Files	2148
List the Core Files Generated by the Crash	2148
Compress the Core Files	2149
Log Software Version Information	2150
Open a Case with JTAC	2151
Display the Crash Stack Traceback and Registration Information	2151
Clear the NVRAM Contents	2156
Check Packet Forwarding Engine Microkernel Core Files	2157
List the Core Files Generated by the Crash	2157
Compress the Core Files	2158
Configure a Primary Path	2159
Ensuring That Secondary Paths Establish When Resources Are Diminished	2161
One-to-One Backup Overview	2163
Configure Link Protection	2164
Configuring and Verifying Link Protection	2166
Configure Link Protection	2166
Verify That Link Protection Is Up	2168
Configure Node-Link Protection	2173
Configuring and Verifying Node-Link Protection	2175
Configure Node-Link Protection	2175
Verify That Node-Link Protection Is Up	2177

Configure IS-IS as the IGP	2185
Enable IS-IS on Routers in Your Network	2187
Configure ISO Addressing	2190
Enable IS-IS on Router Interfaces	2191
Verify That IS-IS Adjacencies Are Established	2193
Verify That IS-IS Adjacencies Are Established	2194
Configure OSPF as the IGP	2196
Enable OSPF on Routers in Your Network	2197
Verify That OSPF Neighbors Are Established	2200
Set Up BGP on Routers in Your Network	2201
Define the Local Autonomous System	2202
Configure BGP Neighbor Connections	2203
Configure a Simple Routing Policy	2205
Verify That BGP Sessions Are Up	2208
Define the Local Autonomous System	2209
Enable MPLS and RSVP	2211
Enable MPLS and RSVP on Routers	2211
Enable MPLS on Transit Router Interfaces	2213
Enable MPLS and RSVP on Routers	2214
Enable MPLS on Transit Router Interfaces	2216
Verifying the MPLS Configuration	2217
Verify the RSVP Protocol	2218
Verify RSVP Interfaces	2220
Verify Protocol Families	2222
Verify the RSVP Protocol	2226
Define a Load-Balancing Policy	2228
Use the traceroute Command to Verify MPLS Labels	2229
Apply the Load-Balancing Policy to the Forwarding Table	2230
Fast Reroute Problem Overview	2232
Problem Establishing a GRE Tunnel Checklist	2257
Troubleshooting GMPLS and GRE Tunnel	2258
Verify Protocol Families	2281

Determining LSP Status | 2285

- Check the Status of the LSP | 2286

- Display Extensive Status About the LSP | 2287

Check the Status of the LSP | 2291

Display Extensive Status About the LSP | 2292

Checking That RSVP Path Messages Are Sent and Received | 2297

Determining the Current RSVP Neighbor State | 2299

Take Appropriate Action | 2300

Examine BGP Routes | 2302

CLI Operational Mode Top-Level Commands | 2303

CLI Keyboard Shortcuts | 2306

Manage Output at the ---(more)--- Prompt | 2307

Working with Problems on Your Network | 2308

Isolating a Broken Network Connection | 2309

Display Junos OS Information | 2310

Display Version Information for Junos OS Packages | 2312

Display the Current Active Router Configuration | 2313

Copy Junos OS to the Router | 2318

Add New Software | 2319

Compare Information Logged Before and After the Upgrade | 2320

Displaying LSP Status Events | 2320

Call Was Cleared by RSVP Event | 2324

Change in Active Path Event | 2324

Clear Call Event | 2325

Deselected as Active Event | 2326

Link Protection Down Event | 2327

Originate Call Event | 2329

ResvTear Received Event | 2329

Session Preempted Event | 2330

Displaying General LSP Error Events | 2331

Admission Control Failure Event | 2333

Explicit Route: Bad Loose Route Event | 2333

Explicit Route: Bad Strict Route Event | 2335

Explicit Route: Format Error Event | 2337

- Explicit Route: Wrong Delivery Event | 2338
- Invalid Destination Address Event | 2340
- Invalid Filter for Policing Event | 2341
- MPLS Graceful Restart: Recovery Failed Event | 2341
- MPLS Label Allocation Failure Event | 2342
- Non-RSVP Capable Router Detected Event | 2343
- No Route Toward Destination Event | 2344
- Unsupported Traffic Class Event | 2345
- CSPF: Computation Result Accepted Event | 2345
- CSPF: Reroute Due to Re-Optimization Event | 2346
- Retry Limit Exceeded Event | 2347
- Log the Software Version Information | 2348
- Log the Hardware Version Information | 2349
- Log the System Boot-Message Information | 2351
- Log the BGP, IS-IS, and OSPF Adjacency Information | 2353
- Back Up the Currently Running and Active File System | 2355
- Reinstall Junos OS | 2355
- Reconfigure Junos OS | 2356
 - Configure Host Names, Domain Names, and IP Addresses | 2356
 - Protecting Network Security by Configuring the Root Password | 2357
 - Check Network Connectivity | 2359
- Configure Host Names, Domain Names, and IP Addresses | 2360
- Check Network Connectivity | 2361
- Automatic Autobandwidth Adjustment Failed Event | 2361
- Configuring Automatic Bandwidth Allocation for LSPs | 2363
 - Configuring Automatic Bandwidth Allocation on LSPs | 2364
 - Configuring the Automatic Bandwidth Allocation Interval | 2365
 - Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth | 2366
 - Configuring the Automatic Bandwidth Adjustment Threshold | 2367
 - Configuring a Limit on Bandwidth Overflow and Underflow Samples | 2368
 - Configuring Passive Bandwidth Utilization Monitoring | 2370
 - Requesting Automatic Bandwidth Allocation Adjustment | 2370
- Displaying DiffServ-Aware Traffic-Engineered LSP Events | 2371
- Unsupported Traffic Class Event | 2372

Traffic Class Value Out of Allowed Range Event | 2373

The Combination of Setup Priority and Traffic Class Is Not One of the Configured TE Classes Event | 2373

RSVP Error, Subcode 7, Signal Type Does Not Match Link Encoding Event | 2374

Unacceptable Label Value Event | 2374

Unsupported Switching Type Event | 2375

Gather Component Alarm Information | 2375

- Display the Current Router Alarms | 2375

- Display Error Messages in the Messages Log File | 2376

- Display Error Messages in the Chassis Process Log File | 2376

Case Study for a CSPF Failure | 2377

- Verify That the LSP Is Established | 2378

- Check the Administrative Group Configuration | 2380

Examining a CSPF Failure | 2385

- Verify the CSPF Failure | 2386

- Examine the CSPF Log File | 2388

- Examine the Traffic Engineering Database | 2390

- Check the Administrative Group Configuration on R5 | 2393

Verify the CSPF Failure | 2396

Examining the Hello Message | 2398

Displaying the Status of IS-IS Adjacencies | 2401

- Verifying Adjacent Routers | 2402

- Examine the Forwarding Table | 2403

Check OSPF on a Stub Router | 2405

Checklist for Verifying the BGP Protocol and Peers | 2406

Verify BGP Peers | 2407

- Verify BGP on an Internal Router | 2409

- Verify BGP on a Border Router | 2412

- Verify Advertised BGP Routes | 2417

- Verify That a Particular BGP Route Is Received on Your Router | 2417

Examine BGP Routes and Route Selection | 2418

- Examine the Local Preference Selection | 2420

- Examine the Multiple Exit Discriminator Route Selection | 2422

- Examine the EBGp over IBGP Selection | 2423

- Examine the IGP Cost Selection | 2425
- Examine the Local Preference Selection | 2426
- Examine the Multiple Exit Discriminator Route Selection | 2427
- Examine the EBGp over IBGP Selection | 2429
- Examine the IGP Cost Selection | 2430
- Examine Routes in the Forwarding Table | 2432
- Ping the Egress Router | 2433
- View the RSVP Log File on Transit Routers | 2434
- Check the RSVP Log File on the Egress Router | 2436
- Determine and Correct the Problem on the Egress Router | 2437
- Check the Routing CPU Memory Usage | 2439
 - Check Overall CPU and Memory Usage | 2439
 - Check Routing Protocol Process (rpd) Memory Usage | 2442
 - Display Tasks | 2446
- Run Snmpwalk from an NMS System to a Juniper Router | 2449
- Configure Trace Operations for SNMP | 2451
- Query a MIB With SNMPGet | 2452
- Check CPU Utilization | 2453
- Check CPU Utilization per Process | 2455
- Retrieve Version Information about Router Software Components | 2459
- Checklist for Displaying Basic Chassis Information | 2460
- Display Basic Chassis Information | 2461
- Maintain a Single Configuration File for Both Routing Engines | 2464
 - Configure the New Group | 2464
 - Apply the New Group | 2466
- Configure the New Group | 2468
- Apply the New Group | 2470
- Viewing Files and Directories on a Device Running Junos OS | 2472
 - Directories on the Device | 2472
 - Listing Files and Directories | 2473
 - Specifying Filenames and URLs | 2476
- Display File Contents | 2477
- Rename a File on a Router | 2478
- Delete a File on a Router | 2479

Check the Time on a Router	2479
Check for Users in Configuration Mode	2480
Check the Commands That Users Are Entering	2481
Configure the Log File for Tracking CLI Commands	2482
Display the Configured Log File	2483
Configure the Log File for Tracking CLI Commands	2484
Check When the Last Configuration Change Occurred	2485
Configure Configuration Change Tracking	2485
Display the Configured Log File	2486
Configure Configuration Change Tracking	2487
Display a Log File	2488
Configure IS-IS-Specific Options	2489
Displaying Detailed IS-IS Protocol Information	2490
Displaying Sent or Received IS-IS Protocol Packets	2493
Analyzing IS-IS Link-State PDUs in Detail	2495
Displaying Detailed IS-IS Protocol Information	2498
Analyzing IS-IS Link-State PDUs in Detail	2501
Configure OSPF-Specific Options	2504
Diagnose OSPF Session Establishment Problems	2504
Analyze OSPF Link-State Advertisement Packets in Detail	2509
Diagnose OSPF Session Establishment Problems	2511
Analyze OSPF Link-State Advertisement Packets in Detail	2516
Chassis	2517
Physical Interface Cards	2518
Routing Engine	2518
Compare Information Logged Before and After the Reinstall	2519
Back Up the New Software	2519
Monitor Hardware Components	2520
Log Software Version Information	2520
Hardware Components	2521
Chassis	2522
Flexible PIC Concentrators	2523

Physical Interface Cards | 2523

Routing Engine | 2523

Power Supplies | 2523

Cooling System | 2524

Configuration Statements

MPLS Configuration Statements | 2527

abstract-hop | 2534

adaptive | 2536

adjust-interval | 2537

adjust-threshold | 2538

adjust-threshold-absolute | 2539

adjust-threshold-activate-bandwidth | 2540

adjust-threshold-overflow-limit | 2541

adjust-threshold-underflow-limit | 2542

admin-down | 2543

admin-group (for Interfaces) | 2544

admin-group (for LSPs) | 2545

admin-group-extended | 2546

admin-groups | 2548

admin-groups-extended | 2549

admin-groups-extended-range | 2551

advertise-mode (MPLS) | 2553

advertisement-hold-time | 2554

allow-fragmentation | 2555

always-mark-connection-protection-tlv | 2556

associate-backup-pe-groups | 2557

associate-lsp | 2558

auto-bandwidth (MPLS Tunnel) | 2559

auto-bandwidth (MPLS Statistics) | 2561

auto-policing | 2562

backup-pe-group | 2563

bandwidth (Fast Reroute, Signaled, and Multiclass LSPs) | 2564

bandwidth (Static LSP) | 2566

bandwidth-model | 2567

bandwidth-percent | 2568

bfd-liveness-detection (Protocols MPLS) | 2569

bfd-liveness-detection (LSP) | 2571

class-of-service (Protocols MPLS) | 2573

compute-options | 2574

compute-profile | 2575

connections (MPLS) | 2577

constituent-list | 2578

container-label-switched-path | 2579

corouted-bidirectional | 2580

corouted-bidirectional-passive | 2581

credibility | 2582

database | 2584

delay (querier) | 2585

delay (responder) | 2586

description (Protocols MPLS) | 2587

description (Protocols Layer 2 VPN) | 2588

deselect-on-bandwidth-failure | 2589

diffserv-te | 2590

disable (Protocols MPLS) | 2591

dual-transport | 2592

dynamic (Source Packet Routing) | 2593

dynamic-tunnels | 2595

egress-protection (MPLS) | 2597

encapsulation-type (Layer 2 VPNs) | 2599

encoding-type | 2601

entropy-label | 2602

entropy-label | 2603

ethernet-vlan (Protocols Link Management) | 2604

ether-pseudowire | 2605

exclude (for Administrative Groups) | 2606

exclude (for Fast Reroute) | 2607

exclude-srlg | 2608

exp | **2609**

expand-loose-hop | **2610**

explicit-null (Protocols MPLS) | **2611**

export (MPLS Traffic engineering database) | **2612**

failure-action (Protocols MPLS) | **2613**

family | **2614**

family mpls | **2615**

fast-reroute (Protocols MPLS) | **2619**

fate-sharing | **2620**

forwarding-rib | **2622**

forwarding-table | **2623**

from (Protocols MPLS) | **2624**

gpip | **2625**

gre (Routing Options) | **2627**

hop-limit | **2628**

import (MPLS Traffic Engineering Database) | **2630**

ip-tunnel-rpf-check | **2632**

ipv4 (Family MPLS) | **2634**

ipv6 (Family MPLS) | **2636**

ip-version (Family MPLS) | **2637**

include-all (for Administrative Groups) | **2638**

include-all (for Fast Reroute) | **2639**

include-any (for Administrative Groups) | **2640**

include-any (for Fast Reroute) | **2641**

ingress (LSP) | **2642**

install (Protocols MPLS) | **2644**

ingress-policy | **2647**

interface (Protocols MPLS) | **2648**

interface (MPLS) | **2649**

inter-domain | **2650**

ip-tunnel-rpf-check | **2651**

ipv6-tunneling | **2652**

label-switched-path (Protocols MPLS) | **2653**

label-switched-path | **2659**

label-switched-path-template (Container LSP) | **2660**

ldp-tunneling | **2661**

least-fill | **2661**

link-protection (Dynamic LSPs) | **2662**

link-protection (Static LSPs) | **2663**

load-balance-label-capability | **2664**

log-updown (Protocols MPLS) | **2665**

longest-match | **2667**

loss (querier) | **2668**

loss (responder) | **2669**

loss-delay (querier) | **2670**

lsp-attributes | **2671**

lsping-channel-type | **2672**

l2vpn | **2673**

maximum-bandwidth (Protocols MPLS) | **2676**

maximum-helper-recovery-time | **2677**

maximum-helper-restart-time (RSVP) | **2678**

maximum-labels | **2679**

minimum-bandwidth-adjust-interval | **2680**

minimum-bandwidth-adjust-threshold-change | **2681**

minimum-bandwidth-adjust-threshold-value | **2682**

metric (Protocols MPLS) | **2683**

minimum-bandwidth | **2684**

monitor-bandwidth | **2685**

most-fill | **2685**

mpls (Protocols) | **2686**

mpls | **2687**

mpls-tp-mode | **2690**

mtu-signaling | **2691**

neighbor (Protocols Layer 2 Circuit) | **2692**

next-hop (Protocols MPLS) | **2694**

no-bfd-triggered-local-repair | **2695**

no-cspf | **2696**

no-decrement-ttl | **2698**

graceful-restart (Enabling Globally) | 2699
 helper-disable (Multiple Protocols) | 2701
 no-install-to-address | 2702
 no-load-balance-label-capability | 2703
 no-mcast-replication | 2704
 no-propagate-ttl | 2705
 no-transit-statistics | 2706
 no-trap | 2707
 node-protection (Static LSP) | 2708
 normalization | 2709
 oam (Protocols MPLS) | 2711
 optimize-adaptive-teardown | 2713
 optimize-aggressive | 2714
 optimize-hold-dead-delay | 2715
 optimize-switchover-delay | 2716
 optimize-timer (Protocols MPLS) | 2717
 p2mp (Protocols MPLS) | 2719
 p2mp-lsp-next-hop | 2720
 path (Protocols MPLS) | 2721
 path | 2723
 path-mtu | 2724
 per-prefix-label | 2725
 performance-monitoring (Protocols MPLS) | 2727
 policing (Protocols MPLS) | 2729
 policing | 2730
 policy-multipath | 2731
 policy-statement | 2733
 pop | 2739
 pop-and-forward (Protocols MPLS) | 2740
 preference (Protocols MPLS) | 2741
 primary (Protocols MPLS) | 2743
 primary | 2744
 priority (Protocols MPLS) | 2745
 protection-revert-time | 2746

push | 2747

random | 2748

record | 2749

remote-interface-switch | 2750

remote-site-id | 2752

retry-limit | 2753

retry-timer | 2754

revert-timer | 2755

revert-timer | 2756

resignal-minimum-bandwidth | 2757

resolution-map | 2758

responder (performance-monitoring) | 2759

rpf-check-policy (Routing Options) | 2760

rsvp-error-hold-time | 2761

sampling (Protocols MPLS) | 2762

sbfd | 2764

secondary (Protocols MPLS) | 2765

secondary | 2767

segment | 2768

segment-list | 2770

select | 2775

signal-bandwidth | 2776

signaling | 2777

site (Layer 2 Circuits) | 2778

site-identifier (Layer 2 Circuits) | 2779

smart-optimize-timer | 2780

soft-preemption (Protocols MPLS) | 2781

source-routing-path | 2782

source-routing-path-template | 2785

source-routing-path-template-map | 2787

splitting-merging | 2788

spring-te (Dynamic Tunnels) | 2790

srgb-label-range | 2791

srlg | 2792

srlg-cost | 2793
srlg-value | 2794
standby | 2795
standby | 2796
static-label-switched-path | 2797
statistics (Protocols MPLS) | 2799
swap | 2801
switch-away-lsps | 2802
switching-type | 2803
sync-active-path-bandwidth | 2804
te-class-matrix | 2806
to | 2808
traceoptions (Protocols MPLS) | 2809
traffic-class (delay) | 2812
traffic-class (loss) | 2814
traffic-class (loss-delay) | 2816
traffic-engineering (Protocols MPLS) | 2818
traffic-engineering | 2820
traffic-engineering (Protocols BGP) | 2821
transit-lsp-association | 2823
ultimate-hop-popping | 2824
vrf-table-label | 2826

RSVP Configuration Statements | 2829

admin-group | 2831
aggregate (Protocols RSVP) | 2832
authentication-key (Protocols RSVP) | 2834
bandwidth (Protocols RSVP) | 2835
bypass (Signaled LSP) | 2837
bypass (Static LSP) | 2839
chained-composite-next-hop | 2840
class-of-service (Protocols RSVP) | 2843
destination-networks | 2844
devices | 2845

disable (Protocols RSVP) | **2846**
dynamic-bidirectional-transport | **2847**
fast-reroute (Protocols RSVP) | **2848**
graceful-deletion-timeout | **2849**
graceful-restart (Protocols RSVP) | **2850**
hello-acknowledgements | **2851**
hello-interval (Protocols RSVP) | **2852**
hop-limit | **2853**
interface (Protocols RSVP) | **2855**
keep-multiplier | **2857**
label-switched-path-template (Multicast) | **2859**
link-protection (RSVP) | **2861**
load-balance (Protocols RSVP) | **2863**
max-bypasses | **2864**
no-local-reversion | **2865**
node-hello | **2867**
no-adjacency-down-notification (Protocols IS-IS) | **2869**
no-authentication-check (Protocols RSVP) | **2870**
no-cspf (Protocols RSVP) | **2871**
no-interface-hello | **2872**
no-neighbor-down-notification | **2873**
no-node-id-subobject | **2874**
no-p2mp-sublsp | **2875**
no-enhanced-frr-bypass (Protocols RSVP) | **2876**
node-link-protection (Protocols MPLS) | **2877**
optimize-timer (Protocols RSVP) | **2878**
path (Protocols RSVP) | **2879**
peer-interface (Protocols RSVP) | **2880**
pop-and-forward (Protocols RSVP) | **2881**
preemption | **2882**
priority (Protocols RSVP) | **2884**
refresh-time | **2885**
reliable | **2886**
rsvp | **2887**

rsvp-te (Routing Options) | 2888
setup-protection | 2889
soft-preemption (Protocols RSVP) | 2890
static-label-switched-path | 2891
subscription | 2893
traceoptions (Protocols RSVP) | 2895
transit | 2898
tunnel-services (RSVP) | 2900
ultimate-hop-popping | 2901
update-threshold | 2903

LDP Configuration Statements | 2905

allow-subnet-mismatch | 2908
authentication-algorithm | 2909
authentication-key (Protocols LDP) | 2912
authentication-key-chain (Protocols LDP) | 2913
auto-targeted-session | 2915
bfd-liveness-detection (Protocols LDP) | 2917
deaggregate | 2919
disable (Protocols LDP) | 2920
dod-request-policy | 2921
downstream-on-demand | 2922
ecmp | 2923
egress-policy | 2924
explicit-null (Protocols LDP) | 2925
export (Protocols LDP) | 2926
failure-action (Protocols LDP) | 2927
fec | 2928
graceful-restart (Protocols LDP) | 2930
hello-interval (Protocols LDP) | 2931
helper-disable (LDP) | 2932
holddown-interval | 2933
hold-time (Protocols LDP) | 2934
ignore-lsp-metrics | 2935

igp-synchronization | 2936

import (Protocols LDP) | 2937

ingress-policy | 2938

interface (Protocols LDP) | 2939

keepalive-interval | 2940

keepalive-timeout | 2941

l2-smart-policy | 2942

label-withdrawal-delay | 2943

ldp | 2944

ldp-synchronization | 2946

log-updown (Protocols LDP) | 2947

make-before-break (LDP) | 2948

mapping-server-entry | 2949

maximum-neighbor-recovery-time | 2950

mldp-inband-signalling (Protocols Multipoint LDP) | 2951

mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain) | 2952

mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain) | 2954

mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain) | 2955

mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute) | 2956

neighbor (Protocols LDP) | 2958

no-forwarding | 2959

oam (Protocols LDP) | 2960

p2mp (Protocols LDP) | 2962

p2mp-ldp-next-hop | 2964

periodic-traceroute | 2965

policing (Protocols LDP) | 2967

policy (Multicast-Only Fast Reroute) | 2968

policy (Protocols Multipoint LDP) | 2971

preference (Protocols LDP) | 2972

prefix-segment (Routing Options) | 2973

prefix-segment-range | 2974

reconnect-time | 2976

recovery-time | 2977

session (Protocols LDP) | 2978

session-group | 2980
session-protection | 2982
source-packet-routing | 2983
stream-protection (Multicast-Only Fast Reroute) | 2984
strict-targeted-hellos | 2985
targeted-hello | 2986
traceoptions (Protocols LDP) | 2987
track-igp-metric | 2990
track-igp-metric (LSP) | 2991
traffic-statistics (Protocols LDP) | 2992
transport-address | 2994
version (BFD) | 2996

CCC and TCC Configuration Statements | 2999

connections (Circuits) | 3000
encapsulation (Logical Interface) | 3002
encapsulation | 3007
interface-switch | 3014
l2circuit-control-passthrough | 3015
lsp-switch | 3016
output-interface (CCC) | 3017
p2mp-receive-switch | 3018
p2mp-transmit-switch | 3019
remote-interface-switch | 3020

GMPLS Configuration Statements | 3021

address (Peer) | 3022
control-channel (Protocols Link Management Peer) | 3023
dead-interval | 3024
disable (GMPLS) | 3025
disable (OSPF) | 3026
export (Protocols BGP) | 3028
hello-dead-interval | 3029
hello-interval (LMP) | 3030

hello-interval (Protocols OSPF) | 3031

import | 3033

instance-type | 3035

interface (Protocols Link Management) | 3038

label-switched-path (Protocols Link Management) | 3039

link-management | 3040

lmp-control-channel | 3042

lmp-protocol | 3043

local-address (Protocols Link Management) | 3044

l2circuit | 3045

passive (Protocols Link Management) | 3047

peer (Protocols LMP) | 3048

peer-interface (Protocols OSPF) | 3049

remote-address (for LMP Control Channel) | 3050

remote-address (for LMP Traffic Engineering) | 3051

remote-id | 3052

retransmission-interval | 3053

retransmit-interval (OSPF) | 3054

retry-limit (Protocols Link Management) | 3056

route-distinguisher | 3057

te-link | 3061

traceoptions (Protocols Link Management) | 3063

transit-delay (OSPF) | 3065

upstream-label | 3067

vrf-target | 3068

PCEP Configuration Statements | 3071

pcep | 3072

delegation-cleanup-timeout | 3073

delegation-priority | 3074

destination-ipv4-address | 3075

destination-port | 3076

label-switched-path-template | 3077

lsp-cleanup-timer | 3078

lsp-external-controller | 3080
max-unknown-messages | 3081
max-unknown-requests | 3082
message-rate-limit | 3083
pce | 3084
pce-group (PCE) | 3088
pce-group (Protocols PCEP) | 3089
pce-type | 3090
querier (performance-monitoring) | 3091
traceoptions (PCE) | 3093
traceoptions (Protocols PCEP) | 3095
update-rate-limit | 3097

Operational Commands

MPLS Operational Commands | 3101

clear mpls lsp | 3103
clear mpls container-lsp | 3105
clear performance-monitoring mpls lsp | 3107
monitor mpls delay rsvp | 3108
monitor mpls loss rsvp | 3114
monitor mpls loss-delay rsvp | 3120
ping mpls bgp | 3125
ping mpls lsp-end-point | 3128
ping mpls l2circuit | 3131
ping mpls l2vpn | 3134
ping mpls l3vpn | 3137
request mpls container-lsp | 3140
request mpls lsp adjust-autobandwidth | 3142
show connections | 3144
show dynamic-tunnels database | 3148
show link-management | 3153
show link-management peer | 3157
show link-management routing | 3160
show link-management statistics | 3164

show link-management te-link | 3167

show mpls abstract-hop-membership | 3170

show mpls admin-groups | 3172

show mpls association | 3174

show mpls call-admission-control | 3176

show mpls container-lsp | 3179

show mpls context-identifier | 3188

show mpls correlation label | 3191

show mpls correlation nexthop-id | 3192

show mpls cspf | 3194

show mpls diffserv-te | 3197

show mpls interface | 3199

show mpls egress-protection | 3201

show mpls interface | 3204

show mpls label usage | 3207

show mpls label usage label-range | 3211

show mpls lsp | 3214

show mpls lsp abstract-computation | 3240

show mpls lsp autobandwidth | 3243

show mpls path | 3246

show mpls srlg | 3248

show mpls static-lsp | 3250

show performance-monitoring mpls lsp | 3254

show route forwarding-table | 3262

show route table | 3273

show ted database | 3329

show ted link | 3341

show ted protocol | 3346

traceroute mpls bgp | 3348

transit (Chained Composite Next Hops) | 3352

RSVP Operational Commands | 3355

clear rsvp session | 3356

clear rsvp statistics | 3358

monitor label-switched-path | 3360

ping mpls rsvp | 3364

show rsvp interface | 3371

show rsvp neighbor | 3379

show rsvp route-session-id | 3385

show rsvp pop-and-forward | 3387

show rsvp session | 3390

show rsvp session | 3405

show rsvp statistics | 3411

show rsvp version | 3419

traceroute mpls rsvp | 3423

LDP Operational Commands | 3429

clear ldp neighbor | 3430

clear ldp session | 3431

clear ldp statistics | 3433

ping mpls ldp | 3434

ping mpls segment routing ospf | 3437

ping mpls segment routing isis | 3440

show ldp database | 3443

show ldp fec-filters | 3454

show ldp interface | 3456

show ldp neighbor | 3458

show ldp overview | 3461

show ldp p2mp tunnel | 3466

show ldp path | 3467

show ldp route | 3469

show ldp session | 3480

show ldp statistics | 3488

show ldp traffic-statistics | 3494

show security keychain | 3499

traceroute mpls ldp | 3502

traceroute mpls segment-routing ospf | 3507

traceroute mpls segment-routing isis | 3511

CCC and TCC Operational Commands | 3515

show connections | **3516**

show route ccc | **3520**

show route forwarding-table | **3522**

PCEP Operational Commands | 3547

clear path-computation-client statistics | **3548**

request path-computation-client active-pce | **3550**

show isis spring sensor info | **3551**

show path-computation-client active-pce | **3554**

show path-computation-client lsp | **3559**

show path-computation-client statistics | **3565**

show path-computation-client status | **3573**

show path-computation-client traffic-steering | **3576**

show spring-traffic-engineering | **3578**

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | [lxv](#)
- Using the Examples in This Manual | [lxv](#)
- Documentation Conventions | [lxvii](#)
- Documentation Feedback | [lxx](#)
- Requesting Technical Support | [lxx](#)

Use this guide to understand the MPLS technology and MPLS applications functions, and to configure MPLS and other feature modules deploying the MPLS applications.

Documentation and Release Notes

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

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 <https://www.juniper.net/books>.

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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


Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see [CLI Explorer](#).

Documentation Conventions

[Table 1 on page lxviii](#) 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 lxviii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions


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

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

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services 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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- 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 <https://support.juniper.net/support/requesting-support/>.

1

PART

MPLS Overview

[Introduction to MPLS | 3](#)

Introduction to MPLS

IN THIS CHAPTER

- [MPLS Overview | 3](#)
- [MPLS Overview | 8](#)
- [MPLS Overview for ACX Series Universal Metro Routers | 13](#)
- [MPLS for EX Series Switches Overview | 14](#)
- [MPLS Applications | 16](#)
- [Link-Layer Support in MPLS | 16](#)
- [Supported MPLS Scaling Values | 16](#)
- [Supported MPLS Standards | 20](#)
- [IP and MPLS Packets on Aggregated Interfaces | 23](#)
- [BGP Destinations | 23](#)
- [IGP and BGP Destinations | 25](#)
- [Install Prefix IGP Overview | 26](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches | 38](#)
- [Understanding MPLS and Path Protection on EX Series Switches | 43](#)
- [MPLS Limitations on QFX Series and EX4600 Switches | 43](#)
- [MPLS Configuration Overview | 48](#)
- [MPLS Configuration Guidelines | 49](#)
- [TTL Processing on Incoming MPLS Packets | 51](#)

MPLS Overview

IN THIS SECTION

- [Why Use MPLS? | 4](#)
- [Why Not Use MPLS? | 5](#)

- [How Do I Configure MPLS? | 5](#)
- [What Does the MPLS Protocol Do? | 6](#)
- [How Does MPLS Interface to Other Protocols? | 7](#)
- [If I Have Used Cisco MPLS, What Do I Need to Know? | 7](#)

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?

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 and Egress Router Addresses for LSPs” on page 514](#) and [“Configuring the Connection Between Ingress and Egress Routers” on page 522](#).

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) 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
tunnel	LSP

What Cisco Calls:	Juniper Calls:
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 | 26](#)

[Understanding MPLS Components for QFX Series and EX4600 Switches | 38](#)

[Understanding MPLS Label Operations | 422](#)

[Understanding CoS MPLS EXP Classifiers and Rewrite Rules | 1345](#)

MPLS Applications User Guide

MPLS Overview

IN THIS SECTION

- [Label Switching | 9](#)
- [Label-Switched Paths | 9](#)
- [Label-Switching Routers | 10](#)
- [Labels | 11](#)
- [Label Operations | 11](#)
- [Penultimate Hop Popping | 12](#)
- [LSP Establishment | 12](#)

Multiprotocol Label Switching (MPLS) is a method for engineering traffic patterns by assigning short labels to network packets that describe how to forward them through the network. MPLS is independent of routing tables or any routing protocol and can be used for unicast packets.

The MPLS framework supports traffic engineering and the creation of virtual private networks (VPNs). Traffic is engineered (controlled) primarily by the use of signaling protocols to establish label-switched paths (LSPs). VPN support includes Layer 2 and Layer 3 VPNs and Layer 2 circuits.

When you enable your device to allow MPLS traffic, the device performs packet-based processing and functions as a standard Junos router.



CAUTION: When packet forwarding mode is changed to MPLS, all flow-based security features are deactivated, and the device performs packet-based processing only. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device. However, MPLS can be enabled in flow-based packet forwarding mode for selected traffic using firewall filters.

Label Switching

In a traditional IP network, packets are transmitted with an IP header that includes a source and destination address. When a router receives such a packet, it examines its forwarding tables for the next-hop address associated with the packet's destination address and forwards the packet to the next-hop location.

In an MPLS network, each packet is encapsulated with an MPLS header. When a router receives the packet, it copies the header as an index into a separate MPLS forwarding table. The MPLS forwarding table consists of pairs of inbound interfaces and path information. Each pair includes forwarding information that the router uses to forward the traffic and modify, when necessary, the MPLS header.

Because the MPLS forwarding table has far fewer entries than the more general forwarding table, the lookup consumes less processing time and processing power. The resultant savings in time and processing are a significant benefit for traffic that uses the network to transit between outside destinations only.


Label-Switched Paths

Label-switched paths (LSPs) are unidirectional routes through a network or autonomous system (AS). In normal IP routing, the packet has no predetermined path. Instead, each router forwards a packet to the next-hop address stored in its forwarding table, based only on the packet's destination address. Each subsequent router then forwards the packet using its own forwarding table.

In contrast, MPLS routers within an AS determine paths through a network through the exchange of MPLS traffic engineering information. Using these paths, the routers direct traffic through the network along an established route. Rather than selecting the next hop along the path as in IP routing, each router is responsible for forwarding the packet to a predetermined next-hop address.

[Figure 1 on page 10](#) shows a typical LSP topology.

Figure 1: Typical LSP Topology



In the topology shown in [Figure 1 on page 10](#), traffic is forwarded from Host C1 to the transit network with standard IP forwarding. When the traffic enters the transit network, it is switched across a preestablished LSP through the network. In this example, an LSP might switch the traffic from Router R4 to Router R2 to Router R1. When the traffic exits the network, it is forwarded to its destination by IP routing protocols.

Label-Switching Routers

Routers that are part of the LSP are label-switching routers (LSRs). Each LSR must be configured with MPLS so that it can interpret MPLS headers and perform the MPLS operations required to pass traffic through the network. An LSP can include four types of LSRs:

- **Inbound router**—The only entry point for traffic into MPLS. Native IPv4 packets are encapsulated into the MPLS protocol by the inbound router. Each LSP can have only one inbound router. Inbound routers are also known as ingress routers.
- **Transit router**—Any router in the middle of an LSP. An individual LSP can contain between 0 and 253 transit routers. Transit routers forward MPLS traffic along the LSP, using only the MPLS header to determine how the packet is routed.
- **Penultimate router**—The second-to-last router in the LSP. The penultimate router in an LSP is responsible for stripping the MPLS header from the packet before forwarding it to the outbound router.
- **Outbound router**—The endpoint for the LSP. The outbound router receives MPLS packets from the penultimate router and performs an IP route lookup. The router then forwards the packet to the next hop of the route. Each LSP can have only one outbound router. Outbound routers are also known as egress routers.

Labels

To forward traffic through an MPLS network, MPLS routers encapsulate packets and assign and manage headers known as *labels*. A label is a 20-bit unsigned integer in the range 0 through 1,048,575. The routers use the labels to index the MPLS forwarding tables that determine how packets are routed through the network.

When a network's inbound router receives traffic, it inserts an MPLS label between the IP packet and the appropriate Layer 2 header for the physical link. The label contains an index value that identifies a next-hop address for the particular LSP. When the next-hop transit router receives the packet, it uses the index in the MPLS label to determine the next-hop address for the packet and forwards the packet to the next router in the LSP.

As each packet travels through the transit network, every router along the way performs a lookup on the MPLS label and forwards the packet accordingly. When the outbound router receives a packet, it examines the header to determine that it is the final router in the LSP. The outbound router then removes the MPLS header, performs a regular IP route lookup, and forwards the packet with its IP header to the next-hop address.

Label Operations

Each LSR along an LSP is responsible for examining the MPLS label, determining the LSP next hop, and performing the required label operations. LSRs can perform five label operations:

- **Push**—Adds a new label to the top of the packet. For IPv4 packets arriving at the inbound router, the new label is the first label in the label stack. For MPLS packets with an existing label, this operation adds a label to the stack and sets the stacking bit to 0, indicating that more MPLS labels follow the first.

When it receives the packet, the inbound router performs an IP route lookup on the packet. Because the route lookup yields an LSP next hop, the inbound router performs a label push on the packet, and then forwards the packet to the LSP next hop.

- **Swap**—Replaces the label at the top of the label stack with a new label.

When a transit router receives the packet, it performs an MPLS forwarding table lookup. The lookup yields the LSP next hop and the path index of the link between the transit router and the next router in the LSP.

- **Pop**—Removes the label from the top of the label stack. For IPv4 packets arriving at the penultimate router, the entire MPLS label is removed from the label stack. For MPLS packets with an existing label, this operation removes the top label from the label stack and modifies the stacking bit as necessary—sets it to 1, for example, if only a single label remains in the stack.

If multiple LSPs terminate at the same outbound router, the router performs MPLS label operations for all outbound traffic on the LSPs. To share the operations among multiple routers, most LSPs use penultimate hop popping (PHP).

- **Multiple push**—Adds multiple labels to the top of the label stack. This action is equivalent to performing multiple push operations.

The multiple push operation is used with label stacking, which is beyond the scope of this topic.

- **Swap and push**—Replaces the top label with a new label and then pushes a new label to the top of the stack.

The swap and push operation is used with label stacking, which is beyond the scope of this topic.

Penultimate Hop Popping

Multiple LSPs terminating at a single outbound router load the router with MPLS label operations for all their outbound traffic. Penultimate hop popping (PHP) transfers the operation from the outbound router to penultimate routers.

With PHP, the penultimate router is responsible for popping the MPLS label and forwarding the traffic to the outbound router. The outbound router then performs an IP route lookup and forwards the traffic. For example, if four LSPs terminate at the same outbound router and each has a different penultimate router, label operations are shared across four routers.

LSP Establishment

An MPLS LSP is established by one of two methods: static LSPs and dynamic LSPs.

Static LSPs

Like a static route, a static LSP requires each router along the path to be configured explicitly. You must manually configure the path and its associated label values. Static LSPs require less processing by the LSRs because no signaling protocol is used. However, because paths are statically configured, they cannot adapt to network conditions. Topology changes and network outages can create black holes in the LSP that exist until you manually reconfigure the LSP.

Dynamic LSPs

Dynamic LSPs use signaling protocols to establish themselves and propagate LSP information to other LSRs in the network. You configure the inbound router with LSP information that is transmitted throughout the network when you enable the signaling protocols across the LSRs. Because the LSRs must exchange and process signaling packets and instructions, dynamic LSPs consume more resources than static LSPs. However, dynamic LSPs can avoid the network black holes of static LSPs by detecting topology changes and outages and propagating them throughout the network.

RELATED DOCUMENTATION

Example: Deleting Security Services

[Example: Enabling MPLS | 56](#)

[MPLS Traffic Engineering and Signaling Protocols Overview | 1134](#)

[MPLS VPN Overview | 1416](#)

MPLS Overview for ACX Series Universal Metro Routers

Multiprotocol Label Switching (MPLS) provides a mechanism for engineering network traffic patterns that is independent of routing tables by assigning short labels to network packets, which describe how to forward them through the network. MPLS is independent of any routing protocol and can be used for unicast packets. On the ACX Series routers, the following MPLS features are supported:

- The configuration of a label-switching router (LSR) for processing of label-switched packets and forwarding of packets based on their labels.
- The configuration of an ingress label edge router (LER) where IP packets are encapsulated within MPLS packets and forwarded to the MPLS domain, and as an egress LER where MPLS packets are decapsulated and the IP packets contained within the MPLS packets are forwarded using information in the IP forwarding table. Configuring MPLS on the LER is the same as configuring an LSR.
- Uniform and pipe mode configuration providing different types of visibility in the MPLS network. Uniform mode makes all the nodes that a label-switched path (LSP) traverses visible to nodes outside the LSP tunnel. Uniform mode is the default. Pipe mode makes only the LSP ingress and egress points visible to nodes outside the LSP tunnel. Pipe mode acts like a circuit and must be enabled with the global **no-propagate-ttl** statement at the `[edit protocols mpls]` hierarchy level on each router that is in the path of the LSP. The **no-propagate-ttl** statement disables time-to-live (TTL) propagation at the router level and affects all RSVP-signalled or LDP-signalled LSPs. Only the global configuration of TTL propagation is supported.
- Exception packet handling of IP packets not processed by the normal packet flow through the Packet Forwarding Engine. The following types of exception packet handling are supported:
 - Router alert
 - Time-to-live (TTL) expiry value
 - Virtual circuit connection verification (VCCV)
- LSP hot standby for secondary paths configuration to maintain a path in a hot-standby state enabling swift cut over to the secondary path when downstream routers on the current active path indicate connectivity problems.
- Redundancy for a label-switched path (LSP) path with the configuration of fast reroute.
- Configuration of link protection to ensure that traffic traversing a specific interface from one router to another can continue to reach its destination in the event that this interface fails.

RELATED DOCUMENTATION

MPLS Applications User Guide

[Disabling Normal TTL Decrementing | 562](#)

[Fast Reroute Overview | 477](#)

[Configuring Fast Reroute | 479](#)

[MPLS and Traffic Protection | 476](#)

[Configuring Link Protection on Interfaces Used by LSPs | 871](#)

[Configuring Hot Standby of Secondary Paths for LSPs | 568](#)

MPLS for EX Series Switches Overview

IN THIS SECTION

- [Benefits of MPLS | 15](#)
- [Additional Benefits of MPLS and Traffic Engineering | 15](#)

You can configure Junos OS MPLS on Juniper Networks EX Series Ethernet Switches to increase transport efficiency in the 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.

NOTE: MPLS configurations on EX Series switches are compatible with configurations on other Juniper Networks devices that support MPLS and MPLS-based circuit cross-connect (CCC). MPLS features available on the switches depend upon which switch you are using. For information about the software features on the EX Series switches, see [Feature Explorer](#).

NOTE: MPLS configurations on the switches do not support:

- Q-in-Q tunneling

This topic describes:

Benefits of MPLS

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.

Additional Benefits of MPLS and Traffic Engineering

MPLS is the packet-forwarding component of the Junos OS traffic engineering architecture. Traffic engineering provides the capabilities to do the following:

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

RELATED DOCUMENTATION

| [FAQ: MPLS on EX Series Switches](#)

MPLS Applications

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

This section discusses the following topics:

- [BGP Destinations on page 23](#)
- [IGP and BGP Destinations on page 25](#)
- [Selecting a Forwarding LSP Next Hop on page 509](#)

Link-Layer Support in MPLS

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

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

Supported MPLS Scaling Values

This topic lists the MPLS scaling values supported on the QFX Series and EX4600 switches.

[Table 3 on page 17](#) lists the MPLS scaling values supported on the QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, EX4600, and EX4650 switches.

NOTE: Starting in Junos OS Release 19.1R1, the ingress tunnel scale is enhanced on QFX Series switches. For instance, you can increase the scale from its default 1024 to 2048 on QFX5100 switches. This enhancement optimizes and increases the ingress tunnel scale to address the needs of data center networks either in IP-CLOS or IP over MPLS application spaces.

Table 3: MPLS Scaling Values for QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, EX4600, and EX4650 Switches

Feature	QFX3500	QFX5100, EX4600	QFX5120, EX4650	QFX5110	QFX5200	QFX5210
Maximum number of MPLS labels in a packet's label stack (push, pop, and swap operations)	3 labels (push) 2 labels (pop) 1 label (swap)	3 labels (push) 2 labels (pop) 1 label (swap)	3 labels (push) 2 labels (pop) 1 label (swap)	3 labels (push) 2 labels (pop) 1 label (swap)	3 labels (push) 2 labels (pop) 1 label (swap)	3 labels (push) 2 labels (pop) 1 label (swap)
Maximum number of MPLS labels (provider switches)	4096	16,384	32,768	32,768	16,384	32,768
Maximum number of tunnel initiators (combination of routes and LSPs)	Ingress LSPs: 1024 Transit LSPs: 4000	Ingress LSPs: 1024 Transit LSPs: 16,384	Ingress LSPs: 8192 Transit LSPs: 15000 Egress LSPs: 8192	Ingress LSPs: 5000 Transit LSPs: 16,384 Egress LSPs: 5000	Ingress LSPs: 2048 Transit LSPs: 15,900 Egress LSPs: 2048	Ingress LSPs: 4096 Transit LSPs: 16,384 Egress LSPs: 4096

Table 3: MPLS Scaling Values for QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, EX4600, and EX4650 Switches (continued)

Feature	QFX3500	QFX5100, EX4600	QFX5120, EX4650	QFX5110	QFX5200	QFX5210
Maximum number of unique next hops (Egress provider edge (PE) switches)	1024	1024	8192	5000	2048	4096
Maximum number of MPLS firewall filters	768	1536	1536(ingress) 2046(egress)	6143 (ingress) 1022 (egress)	768 (ingress) 1024 (egress)	768 (ingress) 1024 (egress)
Virtual routing and forwarding (VRF)	1000	1000	4000	3100	2000	2000
Layer 3 hosts	IPv4: 8000	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>
Layer 3 longest prefix match (LPM)	IPv4: 16,000 IPv6: 4000	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>	See <i>Understanding the Unified Forwarding Table.</i>

Table 4 on page 19 lists the MPLS scaling values supported on the QFX10002, QFX10008, and QFX10016 switches.

Table 4: MPLS Scaling Values for QFX10002, QFX10008, and QFX10016 Switches

Feature	QFX10002	QFX10008	QFX10016
Maximum number of MPLS labels in a packet's label stack (push, pop, and swap operations)	5 labels (push)	5 labels (push)	5 labels (push)
	8 labels (pop)	8 labels (pop)	8 labels (pop)
	1 label (swap)	1 label (swap)	1 labels (swap)
Maximum number of MPLS labels (provider switches); Junos OS limit	128,000	80,000	80,000
Maximum number of tunnel initiators (combination of routes and LSPs); Junos OS limit	Ingress LSPs: 4,096	Ingress LSPs: 4,096	Ingress LSPs: 4,096
	Transit LSPs: 4,096	Transit LSPs: 4,096	Transit LSPs: 4,096
	Egress LSPs: 4,096	Egress LSPs: 4,096	Egress LSPs: 4,096
Maximum number of MPLS firewall filters	8000 (ingress)	8000 (ingress)	8000 (ingress)
	8000 (egress)	8000 (egress)	8000 (egress)
Virtual routing and forwarding (VRF)	4000	4000	4000
Layer 3 hosts	Not applicable	Not applicable	Not applicable
Layer 3 longest prefix match (LPM)	Not applicable	Not applicable	Not applicable

RELATED DOCUMENTATION

[MPLS Configuration Guidelines | 49](#)
[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)
[MPLS Limitations on QFX Series and EX4600 Switches | 43](#)

Supported MPLS Standards

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

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

Only E-LSPs are supported.

- RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*
- RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol*
- RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

Node protection in facility backup is not supported.

- RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*
- RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
- RFC 4385, *Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN.*

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

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

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

- MPLS-TP OAM can send and receive packets with GAL and G-Ach, without IP encapsulation.

- Two unidirectional RSVP LSPs between a pair of routers can be associated with each other to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages. A single Bidirectional Forwarding Detection (BFD) session is established for the associated bidirectional LSP.
- RFC 5712, *MPLS Traffic Engineering Soft Preemption*
- RFC 5718, *An In-Band Data Communication Network For the MPLS Transport Profile*
- RFC 5860, *Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks*
- RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*
- RFC 5921, *A Framework for MPLS in Transport Networks*
- RFC 5950, *Network Management Framework for MPLS-based Transport Networks*
- RFC 5951, *Network Management Requirements for MPLS-based Transport Networks*
- RFC 5960, *MPLS Transport Profile Data Plane Architecture*
- RFC 6215, *MPLS Transport Profile User-to-Network and Network-to-Network Interfaces*
- RFC 6291, *Guidelines for the Use of the "OAM" Acronym in the IETF.*
- RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*
- RFC 6371, *Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks.*
- RFC 6372, *MPLS Transport Profile (MPLS-TP) Survivability Framework*
- RFC 6373, *MPLS-TP Control Plane Framework*
- RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
Only Point-to-Multipoint LSPs are supported.
- RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
- RFC 6425, *Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping*
- RFC 6426, *MPLS On-Demand Connectivity Verification and Route Tracing*
- RFC 6428, *Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile*
- RFC 6510, *Resource Reservation Protocol (RSVP) Message Formats for Label Switched Path (LSP) Attributes Objects*
- RFC 7746, *Label Switched Path (LSP) Self-Ping*
- Internet draft draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt, *Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs*

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

- RFC 2547, *BGP/MPLS VPNs*
- RFC 2702, *Requirements for Traffic Engineering Over MPLS*
- RFC 2917, *A Core MPLS IP VPN Architecture*
- RFC 3063, *MPLS Loop Prevention Mechanism*
- RFC 3208, *PGM Reliable Transport Protocol Specification*

Only the network element is supported.

- RFC 3469, *Framework for Multi-Protocol Label Switching (MPLS)-based Recovery*
- RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*
- Internet draft draft-martini-l2circuit-encap-mpls-11.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*

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

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

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

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

RELATED DOCUMENTATION

[Supported GMPLS Standards | 1356](#)

[Supported LDP Standards | 923](#)

Supported PCEP Standards

[Supported RSVP Standards | 816](#)

Accessing Standards Documents on the Internet

IP and MPLS Packets on Aggregated Interfaces

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

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

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

The Junos OS supports aggregated SONET and Ethernet interfaces.

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

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

RELATED DOCUMENTATION

Ethernet Interfaces User Guide for Routing Devices

Configuring Aggregated SONET/SDH Interfaces

BGP Destinations

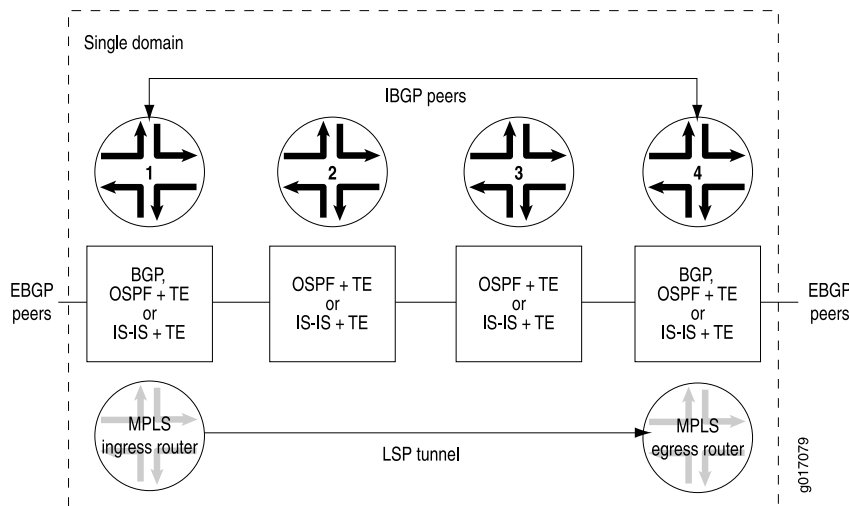
You can configure MPLS to control the paths that traffic takes to destinations outside an AS.

Both IBGP and EBGP take advantage of the LSP host routes without requiring extra configuration. BGP compares the BGP next-hop address with the LSP host route. If a match is found, the packets for the BGP route are label-switched over the LSP. If multiple BGP routes share the same next-hop address, all the BGP routes are mapped to the same LSP route, regardless of which BGP peer the routes are learned from. If the BGP next-hop address does not match an LSP host route, BGP routes continue to be forwarded

based on the IGP routes within the routing domain. In general, when both an LSP route and an IGP route exist for the same BGP next-hop address, the one with the lowest preference is chosen.

Figure 2 on page 24 shows an MPLS topology that illustrates how MPLS and LSPs work. This topology consists of a single domain with four routers. The two routers at the edges of the domain, Router 1 and Router 4, are running EBGp to communicate with peers outside the domain and IBGP to communicate between themselves. For intradomain communication, all four routers are running an IGP. Finally, an LSP tunnel exists from Router 1 to Router 4.

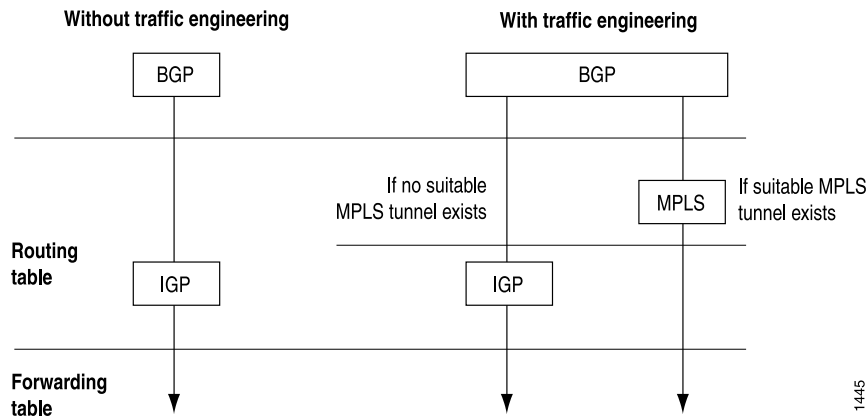
Figure 2: MPLS Application Topology



When BGP on Router 1 receives prefixes from Router 4, it must determine how to reach a BGP next-hop address. Typically, when traffic engineering is not enabled, BGP uses IGP routes to determine how to reach next-hop addresses. (See the left side of Figure 3 on page 25.) However, when traffic engineering is enabled, if the BGP next-hop matches the LSP tunnel endpoint (that is, the MPLS egress router), those prefixes enter the LSP tunnel. (To track these prefixes, look at the **Active Route** field in the **show mpls lsp** command output or at the output of the **show route label-switched-path path-name** command.) If the BGP next hop does not match an LSP tunnel endpoint, those prefixes are sent following the IGP's shortest path.

(See Figure 3 on page 25.)

Figure 3: How BGP Determines How to Reach Next-Hop Addresses



IGP and BGP Destinations

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

When traffic engineering is for BGP destinations only, the MPLS host routes are installed in the `inet.3` routing table (see [Figure 32 on page 474](#)), separate from the routes learned from other routing protocols. Not all `inet.3` routes are downloaded into the forwarding table. Packets directly addressed to the egress router do not follow the LSP, which prevents routes learned from LSPs from overriding routes learned from IGPs or other sources.

Traffic within a domain, including BGP control traffic between BGP peers, is not affected by LSPs. MPLS affects interdomain traffic only; that is, it affects only those BGP prefixes that are learned from an external domain. MPLS does not disrupt intradomain traffic, so IS-IS or OSPF routes remain undisturbed. If you issue a **ping** or **traceroute** command to any destination within the domain, the **ping** or **traceroute** packets follow the IGP path. However, if you issue a **ping** or **traceroute** command from Router 1 in [Figure 2 on page 24](#) (the LSP ingress router) to a destination outside of the domain, the packets use the LSP tunnel.

When traffic engineering for IGP and BGP destinations is enabled, the MPLS host routes are installed in the `inet.0` table (see [Figure 33 on page 475](#)) and downloaded into the forwarding table. Any traffic destined to the egress router could enter the LSP. In effect, it moves all the routes in `inet.3` into `inet.0`, causing the `inet.3` table to be emptied.

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

Install Prefix IGP Overview

When an RSVP LSP is working, the RSVP protocol installs a set of route entries in one or more of the following tables:

- inet.3
- inet6.3
- inet.0
- inet6.0

The details of routes installed in the corresponding tables are controlled by configuration. The routes that may be installed for the LSP by the RSVP protocol are the destination address of the LSP and any number of prefixes configured with the **protocols mpls label-switched-patch *name* install** configuration option.

When installing RSVP protocol routes for an LSP, the resulting metric for the routes can either be a (per-LSP) statically configured value or it can follow the IGP's metric for the LSP's destination address. Before Junos OS release 18.4R1, all RSVP protocol routes would use the same metric. If the LSP had a statically configured metric, then the destination address and all install prefixes would use the statically configured metric. If the LSP used the default behavior of following the IGP metric, then all the install prefixes would use the IGP metric of the LSP's destination address.

From Junos OS release 18.4R1 onwards, a configuration option allows the install prefixes to follow the metric of their corresponding IGP prefix so that the various RSVP protocol routes installed for the LSP can now each have their individual metric value.

RELATED DOCUMENTATION

[track-igp-metric \(LSP\) | 2991](#)

MPLS Feature Support on QFX Series and EX4600 Switches

IN THIS SECTION

- [Supported Features | 27](#)

This topic describes the MPLS features that are supported on the QFX Series, EX4600, EX4650 switches. Be sure to check for any exceptions to this support in [“MPLS Limitations on QFX Series and EX4600 Switches” on page 43](#). Configuring unsupported statements on the switch does not affect its operation.

NOTE: EX4600 and EX4650 switches use the same chipset as QFX5100 switches—this is why EX Series switches are included here along with QFX Series switches. Other EX Series switches also support MPLS but with a different feature set.

Supported Features

The tables in this section lists the MPLS features that are supported on the QFX Series, EX4600, EX4650 switches, and the Junos OS release in which they were introduced. [Table 5 on page 27](#) lists the features for QFX10000 switches. [Table 6 on page 30](#) lists the features for QFX3500, QFX5100, QFX5120, QFX5110, QFX5200, QFX5210 switches. [Table 7 on page 35](#) lists the features for EX4600 and EX4650 switches.

Table 5: QFX10000 MPLS Features

Feature	QFX10002	QFX10008	QFX10016
QFX10000 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-switching router (LSR)	15.1X53-D10	15.1X53-D30	15.1X53-D60
BGP MPLS Ethernet VPN (EVPN)	17.4R1	17.4R1	17.4R1
BGP route reflectors	15.1X53-D10	15.1X53-D30	15.1X53-D60
Automatic bandwidth and dynamic label-switched path (LSP) count sizing	15.1X53-D60	15.1X53-D60, 17.2R1	15.1X53-D60, 17.2R1
BGP labeled unicast	15.1X53-D10	15.1X53-D30	15.1X53-D60
BGP link state distribution	17.1R1	17.1R1	17.1R1
Carrier-of-carriers and interprovider Layer 3 VPNs	17.1R1	17.1R1	17.1R1
Entropy labels	17.2R1	17.2R1	17.2R1

Table 5: QFX10000 MPLS Features (continued)

Feature	QFX10002	QFX10008	QFX10016
Ethernet-over-MPLS (L2 circuit)	15.1X53-D60	15.1X53-D60	15.1X53-D60
Fast reroute, one-to-one local protection and many-to-one local protection	15.1X53-D10	15.1X53-D30	15.1X53-D60
Fast reroute using detours and secondary LSP	15.1X53-D10	15.1X53-D30	15.1X53-D60
Flexible Ethernet services	17.3R1	17.3R1	17.3R1
Firewall filters	15.1X53-D30	15.1X53-D30	15.1X53-D60
RSVP graceful restart for OSPF	15.1X53-D10	15.1X53-D30	15.1X53-D60
IP-over-MPLS LSPs, both static and dynamic links	15.1X53-D10	15.1X53-D30	15.1X53-D60
IPv6 tunneling over an IPv4 network (6PE)	15.1X53-D10	15.1X53-D30	15.1X53-D60
LDP tunneling over RSVP	15.1X53-D10	15.1X53-D30	15.1X53-D60
L2 Circuit on aggregated interfaces	17.3R1	17.3R1	17.3R1
L3VPNs for both IPv4 and IPv6	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
MPLS over UDP	18.3R1	18.3R1	18.3R1
MTU signaling in RSVP	15.1X53-D10	15.1X53-D30	15.1X53-D60
Operation, Administration, and Maintenance (OAM) including ping, traceroute and Bidirectional Forwarding Detection (BFD)	15.1X53-D10	15.1X53-D30	15.1X53-D60
OSPF TE	15.1X53-D10	15.1X53-D30	15.1X53-D60

Table 5: QFX10000 MPLS Features (continued)

Feature	QFX10002	QFX10008	QFX10016
OSPFv2 as an interior gateway protocol (IGP)	15.1X53-D10	15.1X53-D30	15.1X53-D60
Path Computation Element Protocol for RSVP-TE	16.3R1	16.3R1	16.3R1
Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)	15.1X53-D60 (supported only on network-to-network (NNI) interfaces)	15.1X53-D60 (supported only on NNI interfaces)	15.1X53-D60 (supported only on NNI interfaces)
RSVP support, including bandwidth allocation and traffic engineering	15.1X53-D10	15.1X53-D30	15.1X53-D60
RSVP fast reroute (FRR), including link-protection, node-link-protection, fast reroute using detours, and secondary LSP	15.1X53-D10	15.1X53-D30	15.1X53-D60
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 extensions (OSPF-TE, IS-IS-TE)	15.1X53-D10	15.1X53-D30	15.1X53-D60
Traffic engineering (TE) Automatic bandwidth allocation and RSVP bandwidth Dynamic bandwidth management using ingress LSP splitting and merging	15.1X53-D10	15.1X53-D30	15.1X53-D60
Virtual routing and forwarding (VRF) label support	15.1X53-D10	15.1X53-D30	15.1X53-D60

Table 6: QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210 MPLS Features

Feature	QFX3500	QFX5100	QFX5110	QFX5120	QFX5200	QFX5210
QFX Series standalone switches as MPLS provider edge (PE) switches or provider switches	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Label edge router (LER)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Label-switching router (LSR)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Automatic bandwidth allocation on LSPs	Not supported	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
BGP labeled unicast	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
BGP link state distribution	Not supported	17.1R1	17.1R1	18.3R1	17.1R1	18.1R1
BGP route reflector	15.1X53-D10	15.1X53-D30	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Carrier-to-carrier and interprovider BGP Layer 3 VPNs	14.1X53-D15	14.1X53-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1

Table 6: QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210 MPLS Features (continued)

Feature	QFX3500	QFX5100	QFX5110	QFX5120	QFX5200	QFX5210
Class of service (CoS or QoS) for MPLS traffic	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Dynamic label-switched path (LSP) count sizing: TE++	Not supported	17.2R1 VC/VCF 17.2R1	17.2R1 VC/VCF 17.2R1	18.3R1	17.2R1	18.1R1
Equal-cost multipath (ECMP) at LSRs: <ul style="list-style-type: none"> • SWAP • PHP • L3VPN • L2 Circuit 	Not supported	14.1X53-D35 (Supported only on label stack. Not supported on flow label, entropy label, or ECMP label)	15.1X53-D210 (Supported only on label stack. Not supported on flow label, entropy label, or ECMP label)	18.3R1 (Supported only on label stack. Not supported on flow label, entropy label, or ECMP label)	15.1X53-D30	18.1R1
Entropy labels	Not supported	Not supported	Not supported	Not supported	Not supported	18.1R1
Bandwidth over MPLS (L2 Circuit)	14.1X53-D10	14.1X53-D10 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Fast reroute (FRR), one-to-one local protection and many-to-one local protection	14.1X53-D10	14.1X53-D10	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
FRR using detours and secondary LSP	Not supported	Not supported	Not supported	Not supported	Not supported	18.1R1

Table 6: QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210 MPLS Features (continued)

Feature	QFX3500	QFX5100	QFX5110	QFX5120	QFX5200	QFX5210
Firewall filters	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Flow-aware transport of pseudowires (FAT) flow labels	Not supported	Not supported	Not supported	Not supported	Not supported	18.1R1
RSVP graceful restart for OSPF	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Traffic engineering extensions (OSPF-TE, IS-IS-TE)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
IP-over-MPLS LSPs, both static and dynamic links	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
IPv6 tunneling over an MPLS IPv4 network (6PE)	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
IPv6 over an MPLS core network	Not supported	Not supported	Not supported	Not supported	Not supported	Not supported
LDP tunneling over RSVP	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1

Table 6: QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210 MPLS Features (continued)

Feature	QFX3500	QFX5100	QFX5110	QFX5120	QFX5200	QFX5210
Layer 3 VPNs for both IPv4 and IPv6	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Loop-free alternate (LFA)	Not supported	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	18.1R1	18.1R1
MPLS over integrated bridging and routing (IRB) interfaces	Not supported	14.1X53-D40	18.1R1	18.3R1	18.1R1	18.1R1
MTU signaling in RSVP	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Operation, Administration, and Maintenance (OAM) including MPLS ping, traceroute, and BFD	12.3X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
OSPF TE	12.3X50-D10	13.2X51-D15	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
OSPFv2 as an interior gateway protocol	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Path Computation Element Protocol for RSVP-TE	Not supported	17.4R1	17.4R1	18.3R1	17.4R1	18.1R1

Table 6: QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210 MPLS Features (continued)

Feature	QFX3500	QFX5100	QFX5110	QFX5120	QFX5200	QFX5210
Public-facing Ethernet interfaces (core-facing interface)	14.1X53-D10	14.1X53-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
RSVP automatic bandwidth	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
RSVP fast reroute (FRR), including link-protection, node-link-protection, fast reroute using detours, and secondary LSP	14.1X53-D15	14.1X53-D15	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
RSVP-TE extensions (IS-IS and OSPF)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
SNMP MIB support	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Static and dynamic LSPs	12.2X50-D10	13.2X51-D10 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
Traffic engineering (TE) automatic bandwidth allocation on LSPs	13.1X51-D10	13.1X51-D10 VC/VCF (13.2X51-D10)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1

Table 6: QFX3500, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210 MPLS Features (continued)

Feature	QFX3500	QFX5100	QFX5110	QFX5120	QFX5200	QFX5210
Virtual routing and forwarding (VRF) label support	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D210	18.3R1	15.1X53-D30	18.1R1
VRF support in IRB Interfaces in a Layer 3 VPN	Not supported	17.3R1	17.3R1	18.3R1	17.3R1	18.1R1

Table 7: EX4600 and EX4650 MPLS Features

Feature	EX4600	EX4650
EX4600 and EX4650 standalone switches as MPLS provider edge (PE) switches or provider switches	14.1X53-D15	18.3R1
Label edge router (LER)	14.1X53-D15	18.3R1
Label-switching router (LSR)	14.1X53-D15	18.3R1
Automatic bandwidth allocation on LSPs	Not supported	18.3R1
BGP labeled unicast	14.1X53-D15	18.3R1
BGP link state distribution	Not supported	18.3R1
BGP route reflector	14.1X53-D15	18.3R1
Carrier-to-carrier and interprovider BGP Layer 3 VPNs	14.1X53-D15	18.3R1
Class of service (CoS or QoS) for MPLS traffic	14.1X53-D15	18.3R1
Dynamic label-switched path (LSP) count sizing: TE++	Not supported	18.3R1

Table 7: EX4600 and EX4650 MPLS Features (continued)

Feature	EX4600	EX4650
Equal-cost multipath (ECMP) at LSRs: <ul style="list-style-type: none"> • SWAP • PHP • L3VPN • L2 Circuit 	Not supported	18.3R1 (Supported only on label stack. Not supported on flow label, entropy label, or ECMP label)
Entropy labels	Not supported	Not supported
Ethernet-over-MPLS (L2 Circuit)	14.1X53-D15	18.3R1
Fast reroute (FRR), one-to-one local protection and many-to-one local protection	14.1X53-D15	18.3R1
FRR using detours and secondary LSP	Not supported	Not supported
Firewall filters	14.1X53-D15	18.3R1
Flow-aware transport of pseudowires (FAT) flow labels	Not supported	Not supported
RSVP graceful restart for OSPF	13.2X51-D25	18.3R1
Traffic engineering extensions (OSPF-TE, IS-IS-TE)	14.1X53-D15	18.3R1
IP-over-MPLS LSPs, both static and dynamic links	14.1X53-D15	18.3R1
IPv6 tunneling over an MPLS IPv4 network (6PE)	14.1X53-D15	18.3R1
IPv6 over an MPLS core network	Not supported	Not supported
LDP tunneling over RSVP	14.1X53-D15	18.3R1
Layer 3 VPNs for both IPv4 and IPv6	14.1X53-D15	18.3R1
Loop-free alternate (LFA)	Not supported	Not supported

Table 7: EX4600 and EX4650 MPLS Features (continued)

Feature	EX4600	EX4650
MPLS over integrated bridging and routing (IRB) interfaces	Not supported	18.3R1
MTU signaling in RSVP	14.1X53-D15	18.3R1
Operation, Administration, and Maintenance (OAM) including MPLS ping, traceroute, and BFD	14.1X53-D15	18.3R1
OSPF TE	14.1X53-D15	18.3R1
OSPFv2 as an interior gateway protocol	13.2X51-D25	18.3R1
Path Computation Element Protocol for RSVP-TE	Not supported	18.3R1
Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)	14.1X53-D15	18.3R1
RSVP automatic bandwidth	14.1X53-D15	18.3R1
RSVP fast reroute (FRR), including link-protection, node-link-protection, fast reroute using detours, and secondary LSP	14.1X53-D15	18.3R1
RSVP-TE extensions (IS-IS and OSPF)	14.1X53-D15	18.3R1
SNMP MIB support	14.1X53-D15	18.3R1
Static and dynamic LSPs	14.1X53-D15	18.3R1
Traffic engineering (TE) automatic bandwidth allocation on LSPs	14.1X53-D15	18.3R1
Virtual routing and forwarding (VRF) label support	14.1X53-D15	18.3R1
VRF support in IRB Interfaces in a Layer 3 VPN	Not supported	18.3R1

RELATED DOCUMENTATION

[MPLS Configuration Guidelines | 49](#)

[MPLS Limitations on QFX Series and EX4600 Switches | 43](#)

[Supported MPLS Scaling Values | 16](#)

Understanding MPLS Components for QFX Series and EX4600 Switches

IN THIS SECTION

- [Provider Edge Switches | 38](#)
- [Provider Switch | 40](#)
- [Components Required for All Switches in the MPLS Network | 40](#)

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

IN THIS SECTION

- [MPLS Protocol and Label-Switched Paths | 39](#)
- [IP Over MPLS for Customer Edge Interfaces | 39](#)
- [BGP Layer 3 VPN Configuration | 39](#)
- [Routing Instances for Layer 3 VPN | 39](#)
- [Routing Instances for Layer 2 VPN and Layer 3 VPN | 39](#)
- [Ethernet Encapsulation for Layer 2 VPN | 40](#)

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

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 83](#).

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.

Starting in Junos OS Release 15.1, QFX Series devices support the following types of routing instances:

- Layer 2 VPN—To support a Layer 2 VPN (Layer 2 VPNs are not supported on the QFX5100 and EX4600 switches.)
- 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

IN THIS SECTION

- [Interior Gateway Protocol | 40](#)
- [Traffic Engineering | 41](#)
- [MPLS Protocol | 41](#)
- [RSVP | 41](#)
- [Family mpls | 41](#)

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

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 | 26](#)

[Understanding Using MPLS-Based Layer 3 VPNs on Switches | 1420](#)

[Understanding CoS MPLS EXP Classifiers and Rewrite Rules | 1345](#)

[Configuring MPLS on Provider Edge Switches | 83](#)

[Configuring MPLS on Provider Switches | 88](#)

[Configuring Rewrite Rules for MPLS EXP Classifiers | 1349](#)

[Configuring a Global MPLS EXP Classifier | 1352](#)

[Configuring Ethernet over MPLS \(L2 Circuit\) | 1433](#)

MPLS Applications User Guide

Junos OS VPNs Library for Routing Devices

Understanding MPLS and Path Protection on EX Series Switches

Junos OS MPLS for Juniper Networks EX Series Ethernet Switches provides path protection to protect your MPLS network from label switched path (LSP) failures.

By default, an LSP routes itself hop-by-hop from the ingress provider edge switch through the provider switches toward the egress provider edge switch. The LSP generally follows the shortest path as dictated by the local routing table, usually taking the same path as destination-based, best-effort traffic. These paths are “soft” in nature because they automatically reroute themselves whenever a change occurs in a routing table or in the status of a node or link.

Typically, when an LSP fails, the switch immediately upstream from the failure signals the outage to the ingress provider edge switch. The ingress provider edge switch calculates a new path to the egress provider edge switch, establishes the new LSP, and then directs traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress switch might get lost or the new path might take too long to come up, resulting in significant packet drops.

You can configure path protection by configuring primary and secondary paths on the ingress switch. If the primary path fails, the ingress switch immediately reroutes traffic from the failed path to the standby path, eliminating the need for the ingress switch to calculate a new route and signal a new path. For information about configuring standby LSPs, see [“Configuring Path Protection in an MPLS Network \(CLI Procedure\)” on page 145](#).

RELATED DOCUMENTATION

[MPLS for EX Series Switches Overview | 14](#)

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)

[Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)

MPLS Limitations on QFX Series and EX4600 Switches

IN THIS SECTION

- [MPLS Limitations on QFX10000 Switches | 44](#)

- [MPLS Limitations on EX4600, EX4650, QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 Switches | 45](#)

- [MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric Switches | 47](#)
- [MPLS Limitations on QFX3500 Switches | 48](#)

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

MPLS Limitations on QFX10000 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.
- MPLS-over-UDP tunnels are not supported on the following:
 - MPLS TTL propagation
 - IP fragmentation at the tunnel start point
 - CoS rewrite rules and priority propagation for RSVP LSP labels (ingress tunnels only)
 - Plain IPv6
 - Multicast traffic
 - Firewall filters on tunnel start and endpoints
 - CoS tunnel endpoints

NOTE: MPLS-over-UDP tunnels are created only if corresponding RSVP-TE, LDP, or BGP-LU tunnels are not available for the destination route.

MPLS Limitations on EX4600, EX4650, QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 Switches

- MPLS support differs on the various switches. EX4600 switches support only basic MPLS functionality while the QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches support some of the more advanced features. See [“MPLS Feature Support on QFX Series and EX4600 Switches” on page 26](#) for details.
- On a QFX5100 switch, configuring integrated bridging and routing (IRB) interfaces on the MPLS core is implemented on the switch by using TCAM rules. This is the result of a chip limitation on the switch, which only allows for a limited amount of TCAM space. There is 1K TCAM space is allocated for IRB. If multiple IRBs exist, make sure that you have enough available TCAM space on the switch. To check the TCAM space, see [TCAM Filter Space Allocation and Verification in QFX Devices from Junos OS 12.2x50-D20 Onward](#).
- (QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, EX4600) When VLAN bridge encapsulation is enabled on a CE connected interface, the switch drops packets if both flexible Ethernet services and VLAN CCC encapsulations are configured on the same logical interface. Only one can be configured, not both. For example:

set interfaces xe-0/0/18 encapsulation flexible-ethernet-services, or set interfaces xe-0/0/18 encapsulation vlan-ccc.
- Layer 2 circuits on aggregated Ethernet (AE) interfaces are not supported on QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches.
- Layer 2 circuit local switching is not supported on the EX4600, EX4650, and QFX5100 switches.
- The QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches do not depend on the VRF match for loopback filters configured at different routing instances. Loopback filters per routing instance (such as lo0.100, lo0.103, lo0.105) are not supported and may cause unpredictable behavior. We recommend that you only apply the loopback filter (lo0.0) to the master routing instance
- On EX4600 and EX4650 switches, when loopback filters with both accept and deny terms for the same IP address are configured and if RSVP packets have that IP address in either source IP or destination IP, then those RSVP packets will be dropped even if accept terms have higher priority than deny terms. As per design, if the switch receives an RSVP packet with IP OPTION, the packet is copied to the CPU and then the original packet is dropped. Because RSVP packets are marked for drop, the accept term will not process these packets and the deny term will drop the packets.

- On a link-protected, fast reroute Layer 2 circuit, you might see a traffic convergence delay of 200 to 300 milliseconds.
- Layer 2 circuit local switching is not supported on the EX4600, EX4650, and QFX5100 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 Series 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.
- Although fast reroute (FRR) on regular interfaces is supported, the **include-all** and **include-any** options for FRR are not supported. See [“Fast Reroute Overview” on page 477](#).
- FRR is not supported on MPLS over IRB interfaces.
- MPLS-based circuit cross-connects (CCC) are not supported—only circuit-based pseudowires are supported.
- Configuring link aggregation groups (LAGs) on user-to-network interface (UNI) ports for L2 circuits is not supported.
- MTU signaling in RSVP and discovery is supported in the control plane. However, this cannot be enforced in the data plane.
- With L2 circuit-based pseudowires, if multiple equal-cost RSVP LSPs are available to reach an L2 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.
- Firewall filters and policers on **family mpls** are only supported on QFX5100 switches that act as pure label-switching routers (LSRs) in an MPLS network. A pure LSR is a transit router that switches paths solely on the incoming label's instructions. Firewall filters and policers on **family mpls** are not supported on QFX5100 ingress and egress provider edge (PE) switches. This includes switches that perform penultimate hop popping (PHP).
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.
- These are the hardware limitations for EX4600, EX4650, QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches:
 - Push of a maximum of three labels is supported in the MPLS edge switch if label swap is not done.
 - Push of a maximum of two labels is supported in the MPLS edge switch if label swap is done.
 - Pop at line rate is supported for a maximum of two labels.
 - Global label space is supported but interface-specific label space is not supported.
 - MPLS ECMP on PHY node with BOS=1 is not supported for single labels.

- QFX Series 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, EX4600, QFX5100, and QFX5200 switches.
- On EX4600, EX4650, QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches, 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 EX4600, EX4650, QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches:
 - LDP multipoint
 - LDP link protection
 - LDP Bidirectional Forwarding Detection (BFD)
 - LDP Operation Administration and Management (OAM)
 - LDP multicast-only fast reroute (MoFRR)

MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric Switches

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

- Next-hop LSP
- BFD including BFD triggered FRR
- L2 VPN based on BGP (See [RFC 6624](#))
- VPLS
- Extended VLAN CCC
- Pseudowire protection using Ethernet OAM
- Local switching of pseudo-wire
- Pseudowire fault detection based on VCCV
- QFX Series switches with Broadcom chipsets do not support separate next hops for the same label with different S bits (S-0 and S-1). This includes QFX3500, QFX3600, EX4600, QFX5100, and QFX5200 switches.

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 Series 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.
- Although fast reroute is supported, the **include-all** and **include-any** options for fast reroute are not supported. See [“Fast Reroute Overview” on page 477](#) for details.
- 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 the data plane.
- With Layer 2 (L2) circuit-based pseudowires, if multiple equal-cost RSVP label-switched paths (LSPs) are available to reach a L2 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.

RELATED DOCUMENTATION

[MPLS Configuration Guidelines | 49](#)

[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)

[Supported MPLS Scaling Values | 16](#)

MPLS Configuration Overview

When you first install Junos OS on your device, MPLS is disabled by default. You must explicitly configure your device to allow MPLS traffic to pass through. Complete the following steps for all devices in your MPLS network that are running Junos OS.

To enable MPLS:

1. Delete all configured security services from the device. If you do not complete this step, you will get a commit failure. See *Example: Deleting Security Services*.
2. Enable MPLS on the device. See [“Example: Enabling MPLS” on page 56](#).

3. Commit the configuration.
4. Reboot the device.
5. Configure MPLS features such as traffic engineering, VPNs, and VPLS. See:
 - [MPLS Traffic Engineering and Signaling Protocols Overview on page 1134](#)
 - [MPLS VPN Overview on page 1416](#)
 - [CLNS Overview on page 1521](#)
 - [VPLS Overview on page 1546](#)



CAUTION: When packet forwarding mode is changed to MPLS, all flow-based security features are deactivated, and the device performs packet-based processing only. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device. However, MPLS can be enabled in flow-based packet forwarding mode for selected traffic using firewall filters.

RELATED DOCUMENTATION

[MPLS Overview | 8](#)

Example: Deleting Security Services

[Example: Enabling MPLS | 56](#)

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 | 83](#)

[Configuring MPLS on Provider Switches | 88](#)

[Configuring a Global MPLS EXP Classifier | 1352](#)

[Configuring Rewrite Rules for MPLS EXP Classifiers | 1349](#)

[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)

TTL Processing on Incoming MPLS Packets

The flow chart on [Figure 4 on page 52](#) illustrates TTL processing on incoming MPLS packets. On a transit LSR or an egress LER, MPLS pops one or more labels and can push one or more labels. The incoming TTL of the packet is determined by the configured TTL processing tunnel model.

When all of the following conditions are met, the incoming TTL is set to the TTL value found in the immediate inner header:

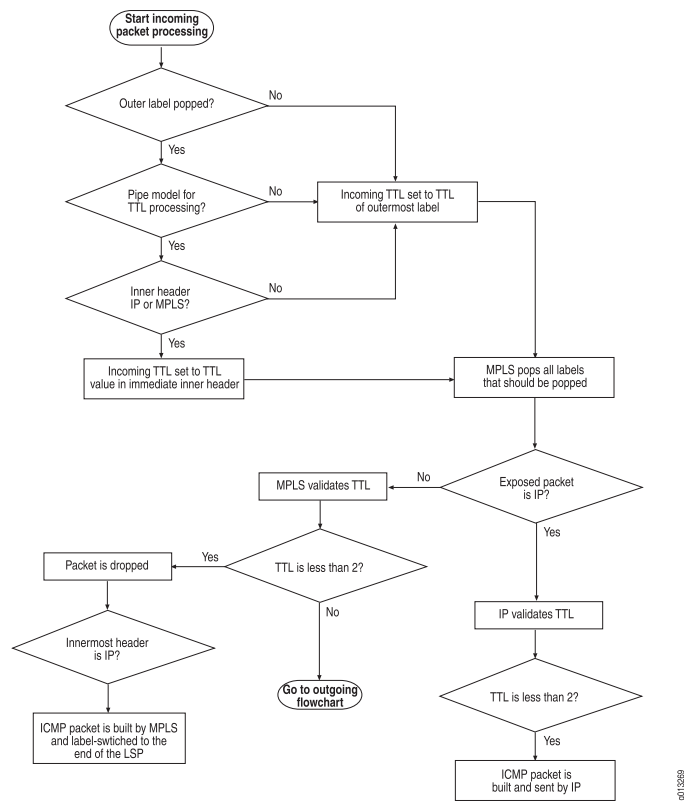
- The outer label is popped as opposed to being swapped
- The TTL processing model is configured to pipe
- The inner header is MPLS or IP

If any of those conditions is not met, then the incoming TTL is set to the TTL value found in the outermost label. In all cases, the TTL values of any further inner labels are ignored.

When an IP packet is exposed after MPLS pops all the labels that should be popped, MPLS passes the packet to IP for further processing, including TTL checking. When the uniform tunnel model for TTL processing is in effect, MPLS sets the TTL value of the IP packet to the incoming TTL value that was just set. In other words, the TTL value is copied from the outermost label to the IP packet. When the pipe model for TTL processing is in effect, the TTL value in the IP header is left unchanged.

If an IP packet is not exposed by the label popping, then MPLS performs the TTL validation. If the incoming TTL is less than 2, the packet is dropped. If innermost packet is IP, an ICMP packet is built and sent. If the TTL does not expire and the packet needs to be sent out, the outgoing TTL is determined by the rules for outgoing MPLS packets.

Figure 4: TTL Processing on Incoming MPLS Packets



RELATED DOCUMENTATION

[Disabling Normal TTL Decrementing | 562](#)
[no-propagate-ttl | 2705](#)

2

PART

Configuring MPLS and Associated Features

Configuring MPLS | **55**

Configuring MPLS on Provider and Provider Edge Devices | **83**

Configuring Bidirectional Forwarding Detection (BFD) for MPLS | **101**

Configuring Firewall Filters, System Log Messages, and SNMP for MPLS | **115**

Configuring Graceful Restart for MPLS | **131**

Configuring IP-Based Filtering for MPLS | **135**

Configuring Link, Node, and Path Protection for MPLS | **143**

Configuring MPLS Load Balancing and Statistics | **239**

Configuring Shared Risk Link Group (SRLG) | **277**

Configuring MPLS Tunnels | **353**

Configuring Policy-Based Multipath Routes | **409**

Configuring MPLS

IN THIS CHAPTER

- [Configuring MPLS | 55](#)
- [Example: Enabling MPLS | 56](#)
- [Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)
- [Verifying That MPLS Is Working Correctly | 79](#)

Configuring MPLS

You must also configure MPLS for a Layer 2 cross-connect to work. The following is a minimal MPLS configuration:

```
[edit]
interfaces {
  interface-name {
    unit logical-unit-number;
  }
}
protocols {
  mpls {
    interface all;
  }
}
```

RELATED DOCUMENTATION

| [Understanding MPLS Label Manager | 425](#)

Example: Enabling MPLS

This example shows how to enable MPLS for packet-based processing. It also shows how to enable the MPLS family and MPLS process on all of the transit interfaces in the network.

NOTE: When MPLS is enabled, all flow-based security features are deactivated and the device performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, IP packets, and IPsec VPNs are unavailable on the device.

Before changing from flow mode to packet mode, you must remove all security policies remaining under flow mode. To prevent management connection loss, you must bind the management interface to zones and enable host-inbound traffic to prevent the device from losing connectivity.

For information about configuring zones, see *Security Policies User Guide for Security Devices*.

Requirements

Before you begin, delete all configured security services. See *Example: Deleting Security Services*.

Overview

The instructions in this topic describe how to enable MPLS on the device. You must enable MPLS on the device before including a device running Junos OS in an MPLS network.

Configuration

CLI Quick Configuration

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

```
set security forwarding-options family mpls mode packet-based
set interfaces ge-1/0/0 unit 0 family mpls
set protocols mpls ge-1/0/0 unit 0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To enable MPLS:

1. Enable MPLS for packet-based processing.

```
[edit security forwarding-options]
user@host# set family mpls mode packet-based
```

2. Enable the MPLS family on each transit interface that you want to include in the MPLS network.

```
[edit interfaces]
user@host# set interfaces ge-1/0/0 unit 0 family mpls
```

3. Enable the MPLS process on all of the transit interfaces in the MPLS network.

```
[edit protocols mpls]
user@host# set interface ge-1/0/0 unit 0
```

Results

From configuration mode, confirm your configuration by entering the **show security forwarding-options** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

NOTE: If you enable MPLS for packet-based processing by using the command **set security forward-option family mpls mode packet**, the mode will not change immediately and the system will display the following messages:

warning: Reboot may required when try reset flow inet mode

warning: Reboot may required when try reset mpls flow mode please check security flow status for detail.

You need to reboot your device for the configuration to take effect.



CAUTION: If you disable MPLS and switch back to using the security services (flow-based processing), the mode will not change immediately and the system will display warning messages instructing you to restart your device. You must reboot your device for the configuration to take effect. This will also result in management sessions being reset and transit traffic getting interrupted.

```
[edit]
user@host# show security forwarding-options
family {
  mpls {
    mode packet-based;
  }
}
```

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

Verification

IN THIS SECTION

- [Verifying MPLS Is Enabled at the Protocols Level | 58](#)
- [Verifying MPLS Is Enabled at the Interfaces Level | 58](#)

Confirm that the configuration is working properly.

Verifying MPLS Is Enabled at the Protocols Level

Purpose

Verify that MPLS is enabled at the protocols level.

Action

From operational mode, enter the **show protocols** command.

Verifying MPLS Is Enabled at the Interfaces Level

Purpose

Verify that MPLS is enabled at the interfaces level.

Action

From operational mode, enter the **show interfaces** command.

RELATED DOCUMENTATION

[MPLS Overview | 8](#)

[MPLS Configuration Overview | 48](#)

Example: Deleting Security Services

Example: Configuring MPLS on EX8200 and EX4500 Switches

IN THIS SECTION

- [Requirements | 60](#)
- [Overview and Topology | 60](#)
- [Configuring the Local PE Switch | 64](#)
- [Configuring the Remote PE Switch | 68](#)
- [Configuring the Provider Switch | 71](#)
- [Verification | 74](#)

You can configure MPLS on switches to increase transport efficiency in your 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.

To implement MPLS on the switches, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch— and at least one provider (transit) switch. You can configure the customer edge (CE) interfaces on the PE switches of the MPLS network as either circuit cross-connect (CCC) or IP (**family inet**) interfaces.

This example shows how to configure an MPLS tunnel using a simple interface as a CCC:

NOTE: This example shows how to configure MPLS using a simple interface as a CCC. For information on configuring a tagged VLAN interface as a CCC, see [“Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\)” on page 1684](#) or [“Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\)” on page 1648](#).

Requirements

This example uses the following hardware and software components:

- Junos OS Release 10.1 or later for switches
- Three EX Series switches

Before you begin configuring MPLS, ensure that you have configured the routing protocol (OSPF or IS-IS) on the core interface and the loopback interface on all the switches. This example includes the configuration of OSPF on all the switches. For information on configuring IS-IS as the routing protocol, see the [Junos OS Routing Protocols Configuration Guide](#).

Overview and Topology

This example includes an ingress or local PE switch, an egress or remote PE switch, and one provider switch. It includes CCCs that tie the customer edge interface of the local PE switch (PE-1) to the customer edge interface of the remote PE switch (PE-2). It also describes how to configure the core interfaces of the PE switches and the provider switch to support the transmission of the MPLS packets. In this example, the core interfaces that connect the local PE switch and the provider switch are individual interfaces, while the core interfaces that connect the remote PE switch and the provider switch are aggregated Ethernet interfaces.

NOTE:

- Core interfaces cannot be tagged VLAN interfaces.
- Core interfaces can be aggregated Ethernet interfaces. This example includes a LAG between the provider switch and the remote PE switch because this type of configuration is another option you can implement. For information on configuring LAGs, see *Configuring Aggregated Ethernet Links (CLI Procedure)*.

[Figure 5 on page 61](#) shows the topology used in this example.

Figure 5: Configuring MPLS on EX Series Switches

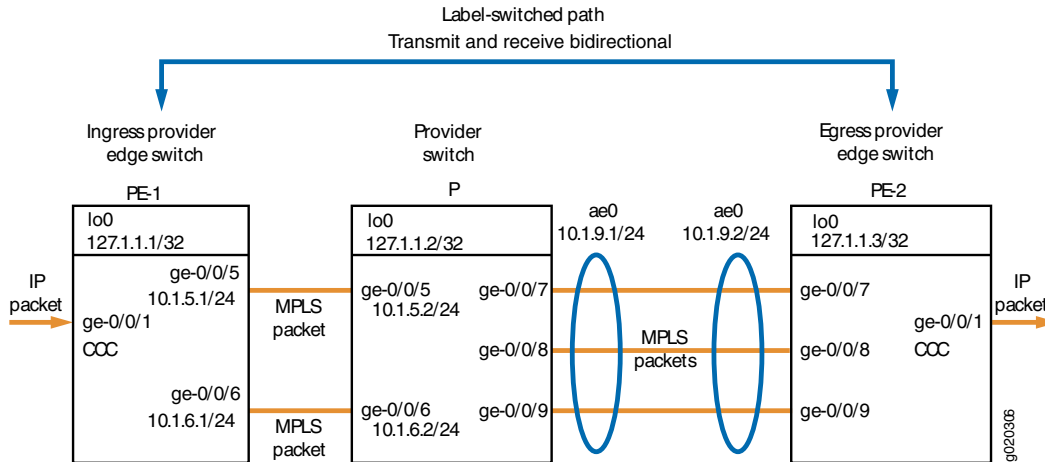


Table 8 on page 61 shows the MPLS configuration components used for the ingress PE switch in this example.

Table 8: Components of the Ingress PE Switch in the Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Local PE switch hardware	EX Series switch	PE-1
Loopback address	lo0 127.1.1.1/32	Identifies PE-1 for interswitch communications.
Routing protocol	ospf traffic-engineering	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol and definition of label-switched path	mpls label-switched-path lsp_to_pe2_ge1 to 127.1.13	Indicates that this PE switch is using the MPLS protocol with the specified LSP to reach the other PE switch (specified by the loopback address). The statement must also specify the core interfaces to be used for MPLS traffic.
RSVP	rsvp	Indicates that this switch is using RSVP. The statement must specify the loopback address and the core interfaces that will be used for the RSVP session.

Table 8: Components of the Ingress PE Switch in the Topology for MPLS with Interface-Based CCC (*continued*)

Property	Settings	Description
Interface family	family inet family mpls family ccc	The logical units of the core interfaces are configured to belong to both family inet and family mpls . The logical unit of the customer edge interface is configured to belong to family ccc .
Customer edge interface	ge-0/0/1	Interface that connects this network to devices outside the network.
Core interfaces	ge-0/0/5.0 and ge-0/0/6.0 with IP addresses 10.1.5.1/24 and 10.1.6.1/24	Interfaces that connect to other switches within the MPLS network.
CCC definition	connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/1.0 transmit-lsp lsp_to_pe2_ge1 receive-lsp lsp_to_pe1_ge1	Associates the circuit cross-connect (CCC), ge-0/0/1 , with the LSPs that have been defined on the local and remote PE switches.

[Table 9 on page 62](#) shows the MPLS configuration components used for the egress PE switch in this example.

Table 9: Components of the Egress PE Switch in the Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Remote PE switch hardware	EX Series switch	PE-2
Loopback address	lo0 127.1.1.3/32	Identifies PE-2 for interswitch communications.
Routing protocol	ospf traffic-engineering	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.

Table 9: Components of the Egress PE Switch in the Topology for MPLS with Interface-Based CCC (*continued*)

Property	Settings	Description
MPLS protocol and definition of label-switched path	mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1	Indicates that this PE switch is using the MPLS protocol with the specified label-switched path (LSP) to reach the other PE switch. The statement must also specify the core interfaces to be used for MPLS traffic.
RSVP	rsvp	Indicates that this switch is using RSVP. The statement must specify the loopback address and the core interfaces that will be used for the RSVP session.
Interface family	family inet family mpls family ccc	The logical unit of the core interface is configured to belong to both family inet and family mpls . The logical unit of the customer edge interface is configured to belong to family ccc .
Customer edge interface	ge-0/0/1	Interface that connects this network to devices outside the network.
Core interface	ae0 with IP address 10.1.9.2/24	Aggregated Ethernet interface on PE-2 that connects to aggregated Ethernet interface ae0 of the provider switch and belongs to family mpls .
CCC definition	connections remote-interface-switch ge-1-to-pe1 interface ge-0/0/1.0 transmit-lsp lsp_to_pe1_ge1; receive-lsp lsp_to_pe2_ge1;	Associates the CCC, ge-0/0/1 , with the LSPs that have been defined on the local and remote PE switches.

Table 10 on page 64 shows the MPLS configuration components used for the provider switch in this example.

Table 10: Components of the Provider Switch in the Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Provider switch hardware	EX Series switch	Transit switch within the MPLS network configuration.
Loopback address	lo0 127.1.1.2/32	Identifies provider switch for interswitch communications.
Routing protocol	ospf traffic-engineering	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol	mpls	Indicates that this switch is using the MPLS protocol. The statement must specify the core interfaces that will be used for MPLS traffic.
RSVP	rsvp	Indicates that this switch is using RSVP. The statement must specify the loopback and the core interfaces that will be used for the RSVP session.
Interface family	family inet family mpls	The logical units for the loopback interface and the core interfaces belong to family inet . The logical units of the core interfaces are also configured to belong to family mpls .
Core interfaces	ge-0/0/5.0 and ge-0/0/6.0 with IP addresses 10.1.5.1/24 and 10.1.6.1/24 and ae0 with IP address 10.1.9.1/24	Interfaces that connect the provider switch (P) to PE-1. Aggregated Ethernet interface on P that connects to aggregated Ethernet interface ae0 of PE-2.

Configuring the Local PE Switch

CLI Quick Configuration

To quickly configure the local ingress PE switch, copy the following commands and paste them into the switch terminal window of PE-1:

```
[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
set protocols mpls interface ge-0/0/5.0
set protocols mpls interface ge-0/0/6.0
set protocols rsvp interface lo0.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-0/0/6.0
set interfaces lo0 unit 0 family inet address 127.1.1.1/32
set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family ccc
set protocols connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/1.0
set protocols connections remote-interface-switch ge-1-to-pe2 transmit-lsp lsp_to_pe2_ge1
set protocols connections remote-interface-switch ge-1-to-pe2 receive-lsp lsp_to_pe1_ge1
```

Step-by-Step Procedure

To configure the local ingress PE switch:

1. Configure OSPF with traffic engineering enabled:

```
[edit protocols]
user@switchPE-1# set ospf traffic-engineering
```

2. Configure OSPF on the loopback address and the core interfaces:

```
[edit protocols]
user@switchPE-1# set ospf area 0.0.0.0 interface lo0.0
user@switchPE-1# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switchPE-1# set ospf area 0.0.0.0 interface ge-0/0/6.0
```

3. Configure MPLS on this PE switch (PE-1) with a label-switched path (LSP) to the other PE switch (PE-2):


```
[edit protocols]
user@switchPE-1# set mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
```

4. Configure MPLS on the core interfaces:

```
[edit protocols]
user@switchPE-1# set mpls interface ge-0/0/5.0
user@switchPE-1# set mpls interface ge-0/0/6.0
```

5. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switchPE-1# set rsvp interface lo0.0
user@switchPE-1# set rsvp interface ge-0/0/5.0
user@switchPE-1# set rsvp interface ge-0/0/6.0
```

6. Configure IP addresses for the loopback interface and the core interfaces:

```
[edit]
user@switchPE-1# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switchPE-1# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switchPE-1# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
```

7. Configure **family mpls** on the logical unit of the core interface addresses:

```
[edit]
user@switchPE-1# set interfaces ge-0/0/5 unit 0 family mpls
user@switchPE-1# set interfaces ge-0/0/6 unit 0 family mpls
```

8. Configure the logical unit of the customer edge interface as a CCC:

```
[edit interfaces ge-0/0/1 unit 0]
-user@PE-1# set family ccc
```

9. Configure the interface-based CCC from PE-1 to PE-2:

NOTE: You can also configure a tagged VLAN interface as a CCC. See [“Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\)”](#) on page 1684 or [“Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\)”](#) on page 1648.

```
[edit protocols]
user@PE-1# set connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/1.0
```



```

user@PE-1# set connections remote-interface-switch ge-1-to-pe2 transmit-lsp lsp_to_pe2_ge1
user@PE-1# set connections remote-interface-switch ge-1-to-pe2 receive-lsp lsp_to_pe1_ge1

```

Results

Display the results of the configuration:

```

user@switchPE-1> show configuration

```

```

interfaces {
  ge-0/0/1 {
    unit 0 {
      family ccc;
    }
  }
  ge-0/0/5 {
    unit 0 {
      family inet {
        address 10.1.5.1/24;
      }
      family mpls;
    }
  }
  ge-0/0/6 {
    unit 0 {
      family inet {
        address 10.1.6.1/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 127.1.1.1/32;
      }
    }
  }
}
protocols {
  rsvp {
    interface lo0.0;
    interface ge-0/0/5.0;
    interface ge-0/0/6.0;
  }
}

```



```

mpls {
    label-switched-path lsp_to_pe2_ge1 {
        to 127.1.1.3;
    }
    interface ge-0/0/5.0;
    interface ge-0/0/6.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/5.0;
        interface ge-0/0/6.0;
    }
}
connections {
    remote-interface-switch ge-1-to-pe2 {
        interface ge-0/0/1.0;
        transmit-lsp lsp_to_pe2_ge1;
        receive-lsp lsp_to_pe1_ge1;
    }
}

```

Configuring the Remote PE Switch

CLI Quick Configuration

To quickly configure the remote PE switch, copy the following commands and paste them into the switch terminal window of PE-2:

```

[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ae0
set protocols mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1
set protocols mpls interface ae0
set protocols rsvp interface lo0.0
set protocols rsvp interface ae0
set interfaces lo0 unit 0 family inet address 127.1.1.3/32
set interfaces ae0 unit 0 family inet address 10.1.9.2/24
set interfaces ae0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family ccc

```



```

set protocols connections remote-interface-switch ge-1-to-pe1 interface ge-0/0/1.0
set protocols connections remote-interface-switch ge-1-to-pe1 transmit-lsp lsp_to_pe1_ge1
set protocols connections remote-interface-switch ge-1-to-pe1 receive-lsp lsp_to_pe2_ge1

```

Step-by-Step Procedure

To configure the remote PE switch (PE-2):

1. Configure OSPF with traffic engineering enabled:

```

[edit protocols]
user@switchPE-2# set ospf traffic-engineering

```

2. Configure OSPF on the loopback interface and the core interface:

```

[edit protocols]
user@switchPE-2# set ospf area 0.0.0.0 interface lo0.0
user@switchPE-2# set ospf area 0.0.0.0 interface ae0

```

3. Configure MPLS on this switch (PE-2) with a label-switched path (LSP) to the other PE switch (PE-1):

```

[edit protocols]
user@switchPE-2# set mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1

```

4. Configure MPLS on the core interface:

```

[edit protocols]
user@switchPE-2# set mpls interface ae0

```

5. Configure RSVP on the loopback interface and the core interface:

```

[edit protocols]
ser@switchPE-2# set rsvp interface lo0.0
user@switchPE-2# set rsvp interface ae0

```

6. Configure IP addresses for the loopback interface and the core interface:

```

[edit]
user@switchPE-2# set interfaces lo0 unit 0 family inet address 127.1.1.3/32
user@switchPE-2# set interfaces ae0 unit 0 family inet address 10.1.9.2/24

```

7. Configure **family mpls** on the logical unit of the core interface:

```

[edit]
user@switchPE-2# set interfaces ae0 unit 0 family mpls

```


8. Configure the logical unit of the customer edge interface as a CCC:

```
[edit interfaces ge-0/0/1 unit 0]
user@PE-2# set family ccc
```

9. Configure the interface-based CCC from PE-2 to PE-1:

```
[edit protocols]
user@PE-2# set connections remote-interface-switch ge-1-to-pe1 interface ge-0/0/1.0
user@PE-2# set connections remote-interface-switch ge-1-to-pe1 transmit-lsp lsp_to_pe1_ge1
user@PE-2# set connections remote-interface-switch ge-1-to-pe1 receive-lsp lsp_to_pe2_ge1
```

Results

Display the results of the configuration:

```
user@switchPE-2> show configuration
```

```
interfaces {
  ge-0/0/1 {
    unit 0 {
      family ccc;
    }
  }
  ae0 {
    unit 0 {
      family inet {
        address 10.1.9.2/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 127.1.1.3/32;
      }
    }
  }
}
protocols {
  rsvp {
    interface lo0.0;
    interface ae0.0;
```



```

    }
    mpls {
        label-switched-path lsp_to_pe1_ge1 {
            to 127.1.1.1;
        }
        interface ae0.0;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface ae0.0;
        }
    }
    connections {
        remote-interface-switch ge-1-to-pe1 {
            interface ge-0/0/1.0;
            transmit-lsp lsp_to_pe1_ge1;
            receive-lsp lsp_to_pe2_ge1;
        }
    }
}

```

Configuring the Provider Switch

CLI Quick Configuration

To quickly configure the provider switch, copy the following commands and paste them into the switch terminal window:

```

[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols ospf area 0.0.0.0 interface ae0
set protocols mpls interface ge-0/0/5.0
set protocols mpls interface ge-0/0/6.0
set protocols mpls interface ae0
set protocols rsvp interface lo0.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-0/0/6.0
set protocols rsvp interface ae0

```



```

set interfaces lo0 unit 0 family inet address 127.1.1.2/32
set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
set interfaces ae0 unit 0 family inet address 10.1.9.1/24
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family mpls
set interfaces ae0 unit 0 family mpls

```

Step-by-Step Procedure

To configure the provider switch:

1. Configure OSPF with traffic engineering enabled:

```

[edit protocols]
user@switchP# set ospf traffic-engineering

```

2. Configure OSPF on the loopback interface and the core interfaces:

```

[edit protocols]
user@switchP# set ospf area 0.0.0.0 interface lo0.0
user@switchP# set ospf area 0.0.0.0 interface ge-0/0/5
user@switchP# set ospf area 0.0.0.0 interface ge-0/0/6
user@switchP# set ospf area 0.0.0.0 interface ae0

```

3. Configure MPLS on the core interfaces on the switch:

```

[edit protocols]
user@switchP# set mpls interface ge-0/0/5
user@switchP# set mpls interface ge-0/0/6
user@switchP# set mpls interface ae0

```

4. Configure RSVP on the loopback interface and the core interfaces:

```

[edit protocols]
user@switchP# set rsvp interface lo0.0
user@switchP# set rsvp interface ge-0/0/5
user@switchP# set rsvp interface ge-0/0/6
user@switchP# set rsvp interface ae0

```

5. Configure IP addresses for the loopback interface and the core interfaces:

```

[edit]

```



```

user@switchP# set interfaces lo0 unit 0 family inet address 127.1.1.2/32
user@switchP# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switchP# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switchP# set interfaces ae0 unit 0 family inet address 10.1.9.1/24

```

6. Configure **family mpls** on the logical unit of the core interface addresses:

```

[edit]
user@switchP# set interfaces ge-0/0/5 unit 0 family mpls
user@switchP# set interfaces ge-0/0/6 unit 0 family mpls
user@switchP# set interfaces ae0 unit 0 family mpls

```

Results

Display the results of the configuration:

```

user@switchP> show configuration

```

```

interfaces {
  ge-0/0/5 {
    unit 0 {
      family inet {
        address 10.1.5.1/24;
      }
      family mpls;
    }
  }
  ge-0/0/6 {
    unit 0 {
      family inet {
        address 10.1.6.1/24;
      }
      family mpls;
    }
  }
  ae0 {
    unit 0 {
      family inet {
        address 10.1.9.1/24;
      }
      family mpls;
    }
  }
}

```



```

}
lo0 {
  unit 0 {
    family inet {
      address 127.1.1.2/32;
    }
  }
}
protocols {
  rsvp {
    interface lo0.0;
    interface ge-0/0/5.0;
    interface ge-0/0/6.0;
    interface ae0.0;
  }
  mpls {
    interface ge-0/0/5.0;
    interface ge-0/0/6.0;
    interface ae0.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface lo0.0;
      interface ge-0/0/5.0;
      interface ge-0/0/6.0;
      interface ae0.0;
    }
  }
}

```

Verification

IN THIS SECTION

- [Verifying the Physical Layer on the Switches | 75](#)
- [Verifying the Routing Protocol | 75](#)
- [Verifying the Core Interfaces Being Used for MPLS Traffic | 76](#)
- [Verifying the Status of the RSVP Sessions | 76](#)
- [Verifying the Assignment of Interfaces for MPLS Label Operations | 77](#)
- [Verifying the Status of the CCC | 78](#)

To confirm that the configuration is working properly, perform these tasks:

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@switchPE-1> show interfaces terse
```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	eth-switch		
ge-0/0/1	up	up			
ge-0/0/1.0	up	up	ccc		
ge-0/0/2	up	up			
ge-0/0/2.0	up	up	eth-switch		
ge-0/0/3	up	up			
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4	up	up			
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5	up	up			
ge-0/0/5.0	up	up	inet	10.1.5.1/24	
			mpls		
ge-0/0/6	up	up			
ge-0/0/6.0	up	up	inet	10.1.6.1/24	
			mpls		

Meaning

The **show interfaces terse** command displays status information about the Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (**Proto** column) shows that interface **ge-0/0/1.0** is configured as a circuit cross-connect. The output for the protocol family of the core interfaces (**ge-0/0/5.0** and **ge-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. Perform this verification task on each of the switches. The state must be **Full**.

Action

```
user@switchPE-1> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
127.1.1.2	ge-0/0/5	Full	10.10.10.10	128	39

Meaning

The **show ospf neighbor** command displays the status of the routing protocol. This 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.

Verifying the Core Interfaces Being Used for MPLS Traffic

Purpose

Verify that the state of the MPLS interface is **Up**. Perform this verification task on each of the switches.

Action

```
user@switchPE-1> 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 the Status of the RSVP Sessions

Purpose

Verify the status of the RSVP sessions. Perform this verification task on each of the switches.

Action

```
user@switchPE-1> show rsvp session
```

```
Ingress RSVP: 1 sessions
```

To	From	State	Rt	Style	Labelin	Labelout	LSPname
----	------	-------	----	-------	---------	----------	---------


```

127.1.13          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**.

Verifying the Assignment of Interfaces for MPLS Label Operations

Purpose

Verify which interface is being used as the beginning of the CCC and which interface is being used to push the MPLS packet to the next hop. Perform this task only on the PE switches.

Action

user@switchPE-1> **show route forwarding-table family mpls**

```

MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                dscd   50    1
0                user   0                recv   49    3
1                user   0                recv   49    3
2                user   0                recv   49    3
299776           user   0                Pop     541    2      ge-0/0/1.0
ge-0/0/1.0 (CCC) user   0 2.0.0.1          Push 299792 540 2      ge-0/0/5.0

```

Meaning

This output shows that the CCC has been set up on interface **ge-0/0/1.0**. The switch receives ingress traffic on **ge-0/0/1.0** and pushes label **299792** onto the packet, which goes out through interface **ge-0/0/5.0**. The output also shows when the switch receives an MPLS packet with label 29976, it pops the label and sends the packet out through interface **ge-0/0/1.0**.

After you have checked the local PE switch, run the same command on the remote PE switch.

Verifying the Status of the CCC

Purpose

Verify the status of the CCC. Perform this task only on the PE switches.

Action

user@switchPE-1> [show connections](#)

```
CCC and TCC connections [Link Monitoring On]
Legend for status (St)                Legend for connection types
UN -- uninitialized                    if-sw:  interface switching
NP -- not present                      rmt-if: remote interface switching
WE -- wrong encapsulation              lsp-sw:  LSP switching
DS -- disabled                        tx-p2mp-sw: transmit P2MP switching
Dn -- down                            rx-p2mp-sw: receive P2MP switching
-> -- only outbound conn is up
<- -- only inbound  conn is up
Up -- operational                      Legend for circuit types
RmtDn -- remote CCC down              intf -- interface
Restart -- restarting                 tlsp -- transmit LSP
                                      rlsp -- receive LSP

Connection/Circuit      Type      St      Time last up      # Up trans
gel-to-pe2              rmt-if    Up      Feb 17 05:00:09    1
  ge-0/0/1.0             intf      Up
  lsp_to_pe1_gel         tlsp      Up
  lsp_to_pe2_gel         rlsp      Up
```

Meaning

The **show connections** command displays the status of the CCC connections. This output verifies that the CCC interface and its associated transmit and receive LSPs are **Up**. After you have checked the local PE switch, run the same command on the remote PE switch.

RELATED DOCUMENTATION

Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect (CLI Procedure) 95
Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) 89
Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure) 99
MPLS for EX Series Switches Overview 14

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 | 79](#)
- 2. [Verifying the Routing Protocol | 80](#)
- 3. [Verifying the Core Interfaces Being Used for the MPLS Traffic | 80](#)
- 4. [Verifying RSVP | 81](#)

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.

Action

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 | 83](#)

[Configuring MPLS on Provider Switches | 88](#)

Configuring MPLS on Provider and Provider Edge Devices

IN THIS CHAPTER

- [Configuring MPLS on Provider Edge Switches | 83](#)
- [Configuring MPLS on Provider Switches | 88](#)
- [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)
- [Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)
- [Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

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 | 84](#)
2. [Configuring the Egress PE Switch | 86](#)

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 | 49](#)

[Configuring MPLS on Provider Switches | 88](#)

[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)

[Understanding MPLS Components for QFX Series and EX4600 Switches | 38](#)

[Understanding CoS MPLS EXP Classifiers and Rewrite Rules | 1345](#)

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
```

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
```

RELATED DOCUMENTATION

[Configuring MPLS on Provider Edge Switches | 83](#)

[MPLS Configuration Guidelines | 49](#)

[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)

[Understanding MPLS Components for QFX Series and EX4600 Switches | 38](#)

[Understanding CoS MPLS EXP Classifiers and Rewrite Rules | 1345](#)

Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)

You can configure MPLS on EX Series switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

To implement MPLS on switches, 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 customer edge (CE) interfaces on the PE switches of the MPLS network by using either IP over MPLS or MPLS over circuit cross-connect (CCC).

The main differences between configuring IP over MPLS and configuring MPLS over CCC are that for IP over MPLS you configure the customer edge interfaces to belong to **family inet** (rather than **family ccc**) and you configure a static route for the label-switched path (LSP). The configuration of the provider switch

is the same regardless of whether you have used IP over MPLS or MPLS over CCC. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)”](#) on page 99.

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 | 90](#)
2. [Configuring the Egress PE Switch | 92](#)

Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 100.100.100.100/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
```

2. Configure OSPF on the loopback and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
```

NOTE: If you want to use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as the core interfaces, replace ge-0/0/5.0 and ge-0/0/6 each with an RVI name (for example, *vlan.logical-interface-number*) or a subinterface name (for example, *interface-name.logical-unit-number*).

RVIs function as logical routers, eliminating the need to have both a switch and a router. Layer 3 subinterfaces allow you to route traffic among multiple VLANs along a single trunk line that connects an EX Series switch to a Layer 2 switch.

3. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

4. Configure RSVP on the loopback interface and the core interfaces:


```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
```

5. Configure MPLS traffic engineering:

```
[edit protocols]
user@switch# set protocols mpls traffic-engineering bgp-igp
```

6. Configure MPLS on the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-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]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
```

8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit]
user@switch# set interfaces ge-2/0/3 unit 0 family inet 121.121.121.1/16
```

9. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface ge-2/0/3.0
```

10. Configure an LSP on the ingress PE switch (100.100.100.100) to send IP packets over MPLS to the egress PE switch (208.208.208.208):

```
[edit protocols mpls]
user@switch# set label-switched-path ip_lspjavae_29 from 100.100.100.100
user@switch# set label-switched-path ip_lspjavae_29 to 208.208.208.208
```

11. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
```



```
user@switch# set label-switched-path ip_lspjavae_29 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:

NOTE: Do not configure a static route if you are using this procedure to configure an MPLS-based Layer 3 VPN.

```
[edit]
```

```
user@switch# set routing-options static route 2.2.2.0/24 next-hop 100.100.100.100
```

```
user@switch# set routing-options 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 for the core interfaces:

```
[edit]
```

```
user@switch# set interfaces lo0 unit 0 family inet address 208.208.208.208/32
```

```
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.20.1/24
```

```
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.21.1/24
```

2. Configure OSPF on the loopback interface (or switch address) and core interfaces:

```
[edit protocols]
```

```
user@switch# set ospf area 0.0.0.0 interface lo0.0
```

```
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
```

```
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
```


NOTE: If you want to use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as the core interfaces, replace ge-0/0/5.0 and ge-0/0/6 each with an RVI name (for example, *vlan.logical-interface-number*) or a subinterface name (for example, *interface-name.logical-unit-number*).

RVIs function as logical routers, eliminating the need to have both a switch and a router. Layer 3 subinterfaces allow you to route traffic among multiple VLANs along a single trunk line that connects an EX Series switch to a Layer 2 switch.

3. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
```

5. Configure MPLS traffic engineering on both BGP and IGP destinations:

```
[edit protocols]
user@switch# set protocols mpls traffic-engineering bgp-igp
```

6. Configure MPLS on the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-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]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
```

8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit]
```



```
user@switch# set interfaces ge-2/0/3 unit 0 family inet address 2.2.2.1/16
```

9. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
```

```
user@switch# set protocols ospf area 0.0.0 interface ge-2/0/3
```

10. Configure an LSP on the egress PE switch (208.208.208.208) to send IP packets over MPLS to the ingress PE switch (100.100.100.100):

```
[edit protocols mpls]
```

```
user@switch# set label-switched-path ip_lsp29_jvae from 208.208.208.208
```

```
user@switch# set label-switched-path ip_lsp29_jvae to 100.100.100.100
```

11. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
```

```
user@switch# set label-switched-path ip_lsp29_jvae 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:

NOTE: Do not configure a static route if you are using this procedure to configure an MPLS-based Layer 3 VPN.

```
[edit]
```

```
user@switch# set routing-options static route 121.121.121.0/24 next-hop 208.208.208.208
```

```
user@switch# set routing-options static route 121.121.121.0/24 resolve
```

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect (CLI Procedure)

Junos OS MPLS for EX8200 and EX4500 switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on switches to increase transport efficiency in your 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 VoIP and other business-critical functions.

This topic describes configuring provider edge (PE) switches in an MPLS network using a circuit cross-connect (CCC). The customer edge interface can be either a simple interface or a tagged VLAN interface.

NOTE: If you are configuring a CCC on a tagged VLAN interface, you do not specify **family ccc**. See [Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN](#) and [Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit](#).

NOTE: If you are going through this procedure in preparation for configuring an MPLS-based Layer 2 VPN, you do not need to configure the association of the label-switched path (LSP) with the customer edge interface. The BGP signaling automates the connections, so manual configuration of the **connections** is not required.

The following guidelines apply to CCC configurations:

- When an interface is configured to belong to **family ccc**, it cannot belong to any other family.
- You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.
- If you are configuring a CCC on a tagged VLAN interface, you must explicitly enable VLAN tagging and specify a VLAN ID. The VLAN ID cannot be configured on logical interface unit **0**. The logical unit number must be **1** or higher. See [Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN](#) and [Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit](#).

This procedure shows how to set up two CCCs:

- If you are configuring a CCC on a simple interface (**ge-0/0/1**), you do not need to enable VLAN tagging or specify a VLAN ID, so you skip those steps.
- If you are configuring a CCC on a tagged VLAN interface (**ge-0/0/2**), include all the steps in this procedure.

To configure a PE switch with a CCC:

1. Configure OSPF (or IS-IS) on the loopback (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
user@switch# set ospf area 0.0.0.0 interface ae0
```

2. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set interfaces ae0 unit 0 family inet address 10.1.9.1/24
```

4. Enable MPLS and define the LSP:

```
[edit protocols]
user@switch# set mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
```

TIP: `lsp_to_pe2_ge1` is the LSP name. You will need to use the specified name again when configuring the CCC.

5. Configure MPLS on the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-0/0/6.0
user@switch# set mpls interface ae0
```

6. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
```



```
user@switch# set rsvp interface ge-0/0/6.0
```

```
user@switch# set rsvp interface ae0
```

7. Configure **family mpls** on the logical units of the core interfaces:

```
[edit]
```

```
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
```

```
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
```

```
user@switch# set interfaces ae0 unit 0 family mpls
```

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

8. If you are configuring a CCC on a tagged VLAN interface, enable VLAN tagging on the customer edge interface **ge-0/0/2** of the local PE switch:

```
[edit interfaces ge-0/0/2]
```

```
user@switch# set vlan-tagging
```

If you are configuring a CCC on a simple interface (**ge-0/0/1**), omit this step.

9. If you are configuring a CCC on a tagged VLAN interface, configure the logical unit of the customer edge interface with a VLAN ID:

```
[edit interfaces ge-0/0/2 unit 1]
```

```
user@switch# set vlan-id 100
```

If you are configuring a CCC on a simple interface (**ge-0/0/1**), omit this step.

10. Configure the logical unit of the customer edge interface to belong to **family ccc**:

- On a simple interface:

```
[edit interfaces ge-0/0/1 unit 0]
```

```
user@switch# set family ccc
```

- On a tagged VLAN interface:

```
[edit interfaces ge-0/0/2 unit 1]
```



```
user@switch# set family ccc
```

11. Associate the CCC interface with two LSPs, one for transmitting MPLS packets and the other for receiving MPLS packets:

NOTE: If you are configuring a Layer 2 VPN, omit this step. The BGP signaling automates the connections, so manual configuration of the **connections** is not required.

- On a simple interface:

```
[edit protocols]
user@switch# set connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/1.0
user@switch# set connections remote-interface-switch ge-1-to-pe2 transmit-lsp lsp_to_pe2_ge1
user@switch# set connections remote-interface-switch ge-1-to-pe2 receive-lsp lsp_to_pe1_ge1
```

- On a tagged VLAN interface:

```
[edit protocols]
user@switch# set connections remote-interface-switch ge-1-to-pe2 interface ge-0/0/2.1
user@switch# set connections remote-interface-switch ge-1-to-pe2 transmit-lsp lsp_to_pe2_ge1
user@switch# set connections remote-interface-switch ge-1-to-pe2 receive-lsp lsp_to_pe1_ge1
```

TIP: The **transmit-lsp** option specifies the LSP name that was configured on PE-1 (the local PE switch) by the **label-switched-path** statement within the **[edit protocols mpls]** hierarchy.

The **receive-lsp** option specifies the LSP name that was configured on PE-2 (the remote PE switch) by the **label-switched-path** statement within the **[edit protocols mpls]** hierarchy.

When you have completed configuring one PE switch, follow the same procedures to configure the other PE switch.

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)

You can configure MPLS on EX8200 and EX4500 switches to increase transport efficiency in your 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 VoIP and other business-critical functions.

To implement MPLS on EX Series switches, you must configure at least one provider switch as a transit switch for the MPLS packets. The configuration of all the provider switches remains the same regardless of whether the provider edge (PE) switches are using circuit cross-connect (CCC) or using MPLS over IP for the customer edge interfaces. Likewise, you do not need to change the configuration of the provider switches if you implement an MPLS-based Layer 2 VPN, Layer 3 VPN, or a Layer 2 circuit configuration.

MPLS requires the configuration of a routing protocol (OSPF or IS-IS) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch. For information on configuring IS-IS as the routing protocol, see [Junos OS Routing Protocols Configuration Guide](#).

To configure the provider switch, complete the following tasks:

1. Enable the routing protocol (OSPF or IS-IS) on the loopback interface and on the core interfaces:

NOTE: You can use the switch address as an alternative to the loopback interface.

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
user@switch# set ospf area 0.0.0.0 interface ae0
```

2. Enable traffic engineering for the routing protocol (traffic engineering must be explicitly enabled for OSPF):

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Enable MPLS within the **protocols** stanza and apply it to the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-0/0/6.0
```



```
user@switch# set mpls interface ae0
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
user@switch# set rsvp interface ae0
```

5. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set interfaces ae0 unit 0 family inet address 10.1.9.2/24
```

6. Configure **family mpls** on the logical units of the core interfaces:

```
[edit]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
user@switch# set interfaces ae0 unit 0 family mpls
```

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

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

[Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) | 1512](#)

Configuring Bidirectional Forwarding Detection (BFD) for MPLS

IN THIS CHAPTER

- [Configuring Bidirectional Forwarding Detection for MPLS \(CLI Procedure\) | 101](#)
- [BFD-Triggered Local Repair for Rapid Convergence | 106](#)
- [Configuring BFD for MPLS IPv4 LSPs | 109](#)

Configuring Bidirectional Forwarding Detection for MPLS (CLI Procedure)

IN THIS SECTION

- [Configuring BFD on Provider Edge and Provider Switches for an LDP-Based LSP | 102](#)
- [Configuring BFD on Provider Edge and Provider Switches for an RSVP-Based LSP | 105](#)

You can configure the Bidirectional Forwarding Detection (BFD) protocol on EX8200 standalone switches and EX8200 Virtual Chassis to detect failures in the MPLS label-switch path (LSP). The BFD protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply from the neighbor after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than those of the failure detection mechanisms for static routes, and thus provide faster detection. These timers are also adaptive. For example, a timer can adapt to a higher value if an adjacency fails, or a neighbor can negotiate a higher value than the one configured.

This topic describes configuring the provider edge (PE) switches and the provider switches to support for LDP-based LSPs and RSVP-based LSPs.

This topic includes:

Configuring BFD on Provider Edge and Provider Switches for an LDP-Based LSP

You can enable BFD for the LDP-based LSPs or RSVP-based LSPs associated with a specific forwarding equivalence class (FEC). Alternatively, you can configure an Operation Administration and Maintenance (OAM) ingress policy to enable BFD on a range of FEC addresses.

Before you configure BFD for an LDP-based based LSP, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\)” on page 89](#).
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

To configure BFD on PE and provider switches:

1. Define an OAM policy:

```
[edit]
user@switch# set protocols ldp oam ingress-policy policy-name
```

2. Specify the FEC on which you want to enable OAM:

```
[edit]
user@switch# set protocols ldp oam fec address
```

3. Specify the minimum transmit and receive interval for the BFD configuration:

NOTE: If you configure the **minimum-interval** statement, you do not need to configure the **minimum-receive-interval** statement or the **minimum-transmit-interval** statement.

```
[edit]
user@switch# set protocols ldp oam bfd-liveness-detection minimum-interval time
```

or

```
[edit]
user@switch# set protocols ldp oam bfd-liveness-detection minimum-receive-interval time
user@switch# set protocols ldp oam bfd-liveness-detection minimum-transmit-interval time
```

4. Specify the detection time multiplier. The negotiated transmit interval multiplied by this value gives the detection time for the receiving system in Asynchronous mode:


```
[edit]
```

```
user@switch# set protocols ldp oam bfd-liveness-detection multiplier multiplier
```

5. Specify the minimum transmit interval (or the minimum receive interval).

```
[edit]
```

```
user@switch# set protocols ldp oam bfd-liveness-detection transmit-interval minimum-interval time
```

6. Specify a threshold for detecting the adaptation of the detection time:

```
[edit]
```

```
user@switch# set protocols ldp oam bfd-liveness-detection detection-time threshold time
```

7. Configure route and next-hop action in the event of a BFD session failure event on the LDP-based LSP:

```
[edit]
```

```
user@switch# set protocols ldp oam bfd-liveness-detection failure-action action
```

NOTE: When a BFD session goes down, you can configure the Junos OS to resignal the LSP path or to simply disable the LSP path. You can configure a standby LSP path to handle traffic while the primary LSP path is unavailable. The switch can automatically recover from LSP failures that can be detected by BFD. By default, if a BFD session fails, the event is simply logged.

8. Specify how long the BFD session must be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

```
[edit]
```

```
user@switch# set protocols ldp oam bfd-liveness-detection holddown-interval time
```

9. Enable tracing of FECs for LDP-based LSPs and specify a source address for sending probes. Then, specify a wait interval, after which to send the probe packet.

```
[edit]
```

```
user@switch# set protocols ldp oam periodic-traceroute source address
```

```
user@switch# set protocols ldp oam periodic-traceroute wait time
```

10. Specify the duration of the LSP ping interval in seconds:

```
[edit]
```



```
user@switch# set protocols ldp oam lsp-ping-interval time
```

11. Specify the action to be taken for the OAM policy:

```
[edit]
```

```
user@switch# set policy-options policy-statement policy-name then accept
```

12. Apply the BFD configurations at the MPLS hierarchy level for the configuration to inherit the statements in the configuration group:

```
[edit]
```

```
user@switch# set apply-groups MPLS
```


Configuring BFD on Provider Edge and Provider Switches for an RSVP-Based LSP

When BFD is configured for an RSVP-based LSP on the ingress switch, it is enabled on the primary path and on all standby secondary paths for that LSP. You can enable BFD for all LSPs on a switch or for specific LSPs. If you configure BFD for a specific LSP, whatever values configured globally for BFD are overridden on that LSP. The BFD sessions originate only at the ingress switch and terminate at the egress switch.

Before you configure BFD for an RSVP-based LSP, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\)” on page 89](#).
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

To configure BFD on PE and provider switches:

1. Specify the minimum transmit and receive interval for the BFD configuration:

NOTE: If you configure the **minimum-interval** statement, you do not need to configure the **minimum-receive-interval** statement or the **minimum-transmit-interval** statement.

```
[edit]
```

```
user@switch# set protocols mpls label-switched-path lsp-name oam bfd-liveness-detection  
minimum-interval time
```

or

```
[edit]
```

```
user@switch# set protocols mpls label-switched-path lsp-name oam bfd-liveness-detection  
minimum-receive-interval time
```

```
user@switch# set protocols mpls label-switched-path lsp-name oam bfd-liveness-detection  
minimum-transmit-interval time
```

2. Specify the detection time multiplier. The negotiated transmit interval multiplied by this value gives the detection time for the receiving system in Asynchronous mode:

```
[edit]
```

```
user@switch# set protocols mpls label-switched-path lsp-name oam bfd-liveness-detection multiplier  
multiplier
```

3. Specify the minimum transmit interval (or the minimum receive interval):


```
[edit]
```

```
user@switch# set protocols mpls label-switched-path lsp-name oam bfd-liveness-detection  
transmit-interval minimum-interval time
```

4. Configure route and next-hop actions in the event of a BFD session failure event on the RSVP-based LSP:

```
[edit]
```

```
user@switch# set protocols mpls label-switched-path lsp-name oam bfd-liveness-detection  
failure-action action
```

NOTE: When a BFD session goes down, you can configure the Junos OS to resignal the LSP path or to simply disable the LSP path. You can configure a standby LSP path to handle traffic while the primary LSP path is unavailable. The switch can automatically recover from LSP failures that can be detected by BFD. By default, if a BFD session fails, the event is simply logged if you do not specifically configure a failure action.

RELATED DOCUMENTATION

| [Example: Configuring MPLS on EX8200 and EX4500 Switches](#) | 59

BFD-Triggered Local Repair for Rapid Convergence

IN THIS SECTION

- [Understanding BFD-Triggered Local Protection](#) | 107

Understanding BFD-Triggered Local Protection

IN THIS SECTION

- Purpose of BFD-Triggered Local Repair | 107
- Configuring BFD-Triggered Local Repair | 108
- Disabling BFD-Triggered Local Repair | 108

The time it takes for a network to converge following a link or node failure can vary dramatically based on a number of factors, including network size, the protocols used, and network design. However, while each particular convergence event is different, the process of convergence is essentially consistent. The failure is detected, the failure is reported (flooded) in the network, an alternate path is found for traffic, and the forwarding plane is updated to pass traffic on a new path.

This overview discusses how Bidirectional Forwarding Detection (BFD)-triggered local repair contributes to a quicker restoration time for rapid convergence in an MPLS network.

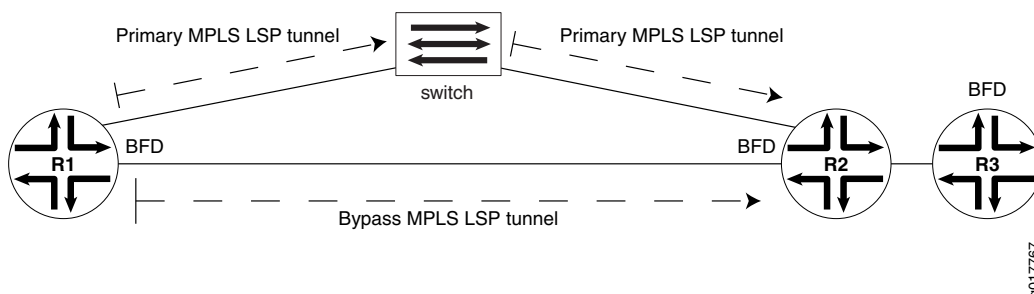
Purpose of BFD-Triggered Local Repair

In Junos OS, general MPLS traffic protection for RSVP-signaled label-switched path (LSP) failures is provided by several complementary mechanisms. These protection mechanisms include local protection (fast reroute, link protection, and node-link protection) and path protection (primary and secondary paths). Local protection in conjunction with path protection can provide minimum packet loss for an LSP, and control the way the LSP is rerouted after a failure. Traditionally, both types of protection rely on fast detection of connectivity failure at the physical level. However, for transmission media without fast physical level detection, Junos OS supports BFD and MPLS ping for fast failure detection.

With links between routers, when a route goes down, the routing protocol process recalculates the next best path. When MPLS fast reroute (FRR) is enabled, ifl messages are flooded to all Flexible PIC Concentrators (FPCs). The edge FPC enables the bypass MPLS LSP tunnel. Lastly, all routes are repaired and sent through the bypass MPLS LSP tunnel. The amount of time it takes to repair all routes is proportional to the number of routes.

This repair scenario becomes more difficult when a switch lies between two links. See [Figure 6 on page 108](#).

Figure 6: Topology with BFD-Triggered Local Repair



When a link goes down at the remote end, the failure is not detected at the local end until the interior gateway protocol (IGP) goes down. To wait for the routing protocol process to recalculate the next best path takes too much time.

With BFD-triggered local repair enabled, the Packet Forwarding Engine completes the repair first, using the bypass MPLS LSP tunnel (that is preconfigured and installed), then informs the routing protocol process to recalculate a new route. By doing this, when the primary MPLS LSP tunnel goes down, the FPC can intermittently and immediately divert traffic to the FPC with the bypass MPLS LSP tunnel.

Using local repair in this way achieves a faster restoration time of less than 50 ms.

Configuring BFD-Triggered Local Repair

BFD-triggered local repair is not configurable, but is part of the default configuration.

BFD-triggered local repair works within the legacy Junos OS features MPLS-FRR, BFD for IGP, and loop-free alternates (LFAs).

Disabling BFD-Triggered Local Repair

By default, BFD-triggered local repair is enabled for all routing interfaces. If desired, you can disable BFD-triggered local repair at the **[edit routing-options]** hierarchy level.

To explicitly disable BFD-triggered local repair:

1. Include the **no-bfd-triggered-local-repair** statement at the **[edit routing-options]** hierarchy level:

```
user@host# set no-bfd-triggered-local-repair
```

2. (Optional) Verify your configuration settings before committing them by using the **show routing-options** command.

```
user@host# run show routing-options
```

Confirm your configuration by issuing the **show routing-options** command.


```
user@host# show routing-options
```

```
...
```

```
no-bfd-triggered-local-repair;
```

```
}
```

NOTE: When you disable this feature, you must also restart routing by including the **graceful-restart** statement for the IGP. For example, for OSPF, this is accomplished by including the **graceful-restart** statement at the **[edit protocols ospf]** hierarchy level.

RELATED DOCUMENTATION

[Fast Reroute Overview | 477](#)

[Configuring BFD for LDP LSPs | 975](#)

[Configuring Link Protection on Interfaces Used by LSPs | 871](#)

[Configuring Fast Reroute | 479](#)

[Configuring Graceful Restart for Point-to-Multipoint LSPs | 724](#)

graceful-restart (Protocols OSPF)

Configuring BFD for MPLS IPv4 LSPs

You can configure Bidirectional Forwarding Detection (BFD) protocol on MPLS IPv4 LSPs as outlined in the Internet draft draft-ietf-bfd-mpls-02.txt, *BFD for MPLS LSPs*. BFD is used as a periodic Operation, Administration, and Maintenance (OAM) feature for LSPs to detect LSP data plane faults. You can configure BFD for LSPs that use either LDP or RSVP as the signaling protocol.

NOTE: BFD for MPLS IPv4 LSP is based on the Routing Engine and is not distributed. As a result, the minimum supported BFD timer interval is (100 ms * 3) per one LSP session, and for scaled LSP sessions, the minimum supported BFD timer interval is (300 ms * 3). As you increase the number of LSP sessions with BFD, you must also increase (scale) the interval timers to support the network.

For Routing Engine switchover instances with nonstop active routing (NSR) support, the minimum supported BFD timer interval is (2.5 seconds * 3).

You can also use the **LSP ping** commands to detect LSP data plane faults. However, BFD has a couple of benefits: it requires less computer processing than **LSP ping** commands and can quickly detect faults in large numbers of LSPs (**LSP ping** commands must be issued for each LSP individually). On the other hand, BFD cannot be used to verify the control plane against the data plane at the egress LSR, which is possible when an **LSP ping** echo request is associated with a forwarding equivalence class (FEC).

The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

Starting from Junos OS Release 13.2R4, 13.3R2, and 14.1, you can set the time interval between LSP ping messages and the number of LSP ping responses, respectively, after which the Bidirectional Forwarding Detection (BFD) session is brought down. To so do, you configure the **lsp-ping-interval** statement and the **lsp-ping-multiplier** statement at the **[edit protocols mpls oam]** hierarchy level.

For configuration instructions for LDP-signaled LSPs, see [“Configuring BFD for LDP LSPs” on page 975](#). For configuration instructions for RSVP-signaled LSPs, see the following section.

Configuring BFD for RSVP-Signaled LSPs

BFD for RSVP supports unicast IPv4 LSPs. When BFD is configured for an RSVP LSP on the ingress router, it is enabled on the primary path and on all standby secondary paths for that LSP. The source IP address for outgoing BFD packets from the egress side of an MPLS BFD session is based on the outgoing interface IP address. You can enable BFD for all LSPs on a router or for specific LSPs. If you configure BFD for a specific LSP, whatever values configured globally for BFD are overridden. The BFD sessions originate only at the ingress router and terminate at the egress router.

An error is logged whenever a BFD session for a path fails. The following example shows how BFD for RSVP LSP log messages might appear:

```
RPD_MPLS_PATH_BFD_UP: MPLS BFD session for path path1 up on LSP R0_to_R3
RPD_MPLS_PATH_BFD_DOWN: MPLS BFD session for path path1 down on LSP R0_to_R3
```

You can configure BFD for all of the RSVP LSPs on the router, a specific LSP, or the primary path of a specific LSP. To configure BFD for RSVP LSPs, include the **oam** and **bfd-liveness-detection** statements.

```
oam {
```



```

bfd-liveness-detection {
  failure-action {
    make-before-break teardown-timeout seconds;
    teardown;
  }
  failure-action teardown;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  minimum-transmit-interval milliseconds;
  multiplier detection-time-multiplier;
}
lsp-ping-interval time-interval;
lsp-ping-multiplier multiplier;
}

```

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit protocols mpls **label-switched-path** *lsp-name* **primary** *path-name*]

The **bfd-liveness-detection** statement includes the following options:

- **minimum-interval**—Specifies the minimum transmit and receive interval.
- **minimum-receive-interval**—Specifies the minimum receive interval. The range is from 1 through 255,000 milliseconds.
- **minimum-transmit-interval**—Specifies the minimum transmit interval. The range is from 1 through 255,000 milliseconds.
- **lsp-ping-multiplier**—Specifies the detection time multiplier. The range is from 1 through 255.

NOTE: To avoid triggering false negatives, configure a BFD fault detection time that is longer than the fast reroute time.

You can also configure the **lsp-ping-interval** option to adjust the time interval between LSP pings. The LSP ping command for RSVP-signaled LSPs is **ping mpls rsvp**. For more information on the **ping mpls rsvp** command, see the [CLI Explorer](#).

Configuring a Failure Action for the BFD Session on an RSVP LSP

When the BFD session for an RSVP LSP goes down, the LSP is torn down and resigaled. Traffic can be switched to a standby LSP, or you can simply tear down the LSP path. Any actions performed are logged.

When a BFD session for an RSVP LSP path goes down, you can configure the Junos OS to resignal the LSP path or to simply disable the LSP path. A standby LSP path could be configured to handle traffic while the primary LSP path is unavailable. The router can automatically recover from LSP failures that can be detected by BFD. By default, if a BFD session fails, the event is simply logged.

To enable the Junos OS to tear down an RSVP LSP path in the event of a BFD event, include the **failure-action** statement:

```
failure-action {
  make-before-break teardown-timeout seconds;
  teardown;
}
```

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

You can configure either the **teardown** or **make-before-break** options:

- **teardown**—Causes the LSP path to be taken down and resigaled immediately.
- **make-before-break**—Causes the Junos OS to attempt to signal a new LSP path before tearing down the old LSP path. You can also configure the **teardown-timeout** option to automatically tear down the LSP after the time period specified if the attempt to resignal the LSP fails within the **teardown-timeout** interval. If you specify a value of 0 for the **teardown-timeout** interval, the LSP is taken down and resigaled immediately (the same behavior as when you configure the **teardown** option).

To configure a failure action for all of the RSVP LSPs, include the **failure-action** statement at the **[edit protocols mpls oam bfd-liveness-detection]** hierarchy level. To configure a failure action for a specific RSVP LSP, include the **failure-action** statement at the **[edit protocols mpls label-switched-path *lsp-name* oam bfd-liveness-detection]** hierarchy level.

To configure a failure action for a specific primary path, include the **failure-action** statement at the **[edit protocols mpls label-switched path *lsp-name* primary *path-name* oam bfd-liveness-detection]** hierarchy level. To configure a failure action for a specific secondary LSP path, include the **failure-action** statement at the **[edit protocols mpls label-switched-path *lsp-name* secondary *path-name* oam bfd-liveness-detection]** hierarchy level.

Release History Table

Release	Description
13.2R4	Starting from Junos OS Release 13.2R4, 13.3R2, and 14.1, you can set the time interval between LSP ping messages and the number of LSP ping responses, respectively, after which the Bidirectional Forwarding Detection (BFD) session is brought down.

Configuring Firewall Filters, System Log Messages, and SNMP for MPLS

IN THIS CHAPTER

- [Configuring MPLS Firewall Filters and Policers on Switches | 115](#)
- [Configuring MPLS Firewall Filters and Policers on Routers | 119](#)
- [System Log Messages and SNMP Traps for MPLS | 128](#)

Configuring MPLS Firewall Filters and Policers on Switches

IN THIS SECTION

- [Configuring an MPLS Firewall Filter | 116](#)
- [Applying an MPLS Firewall Filter to an MPLS Interface | 116](#)
- [Applying an MPLS Firewall Filter to a Loopback Interface | 117](#)
- [Configuring Policers for LSPs | 118](#)

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.

When you configure an MPLS firewall filter, you define the filtering criteria (terms, with match conditions) and an action for the switch to take if the packets match the filtering criteria.

NOTE: You can only configure MPLS filters in the ingress direction. Egress MPLS firewall filters are not supported.

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
```

3. When you are finished, follow the steps below to apply the filter to an interface.

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
```

Applying an MPLS Firewall Filter to a Loopback Interface

To apply an MPLS firewall filter to a loopback interface (lo0):

1. First, specify the packet format by using the *packet-format-match* command. You must restart the PFE every time you configure this command.
2. Configure the firewall filter match conditions and actions as described in [“Configuring an MPLS Firewall Filter” on page 116](#). You must explicitly set the TTL match condition to (**ttl=1**). You can also match packets with other MPLS qualifiers such as **label**, **exp**, and Layer 4 **source port**, and **destination port**.
3. Apply the filter to the loopback interface as an input filter.

```
[edit interfaces]
user@switch# set lo0 unit 0 family mpls filter input ingress-exp-filter
```

4. Review your configuration and issue the **commit** command:

```
[edit interfaces]
user@switch# commit
commit complete
```

The following is an example configuration.

```
set groups lo_mpls_filter interfaces lo0 unit 0 family mpls filter input mpls_lo
set groups lo_mpls_filter firewall family mpls filter mpls_lo term mpls_lo_term from ttl 1
set groups lo_mpls_filter firewall family mpls filter mpls_lo term mpls_lo_term from ip-version ipv4 protocol
  udp source-port 10
set groups lo_mpls_filter firewall family mpls filter mpls_lo term mpls_lo_term from ip-version ipv4 protocol
  udp destination-port 11
set groups lo_mpls_filter firewall family mpls filter mpls_lo term mpls_lo_term then count c1
set groups lo_mpls_filter firewall family mpls filter mpls_lo term mpls_lo_term then accept
```


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 | 26](#)

[Supported MPLS Scaling Values | 16](#)

[Overview of MPLS Firewall Filters on Loopback Interface](#)

[Overview of Policers](#)

Configuring MPLS Firewall Filters and Policers on Routers

IN THIS SECTION

- [Configuring MPLS Firewall Filters | 119](#)
- [Examples: Configuring MPLS Firewall Filters | 120](#)
- [Configuring Policers for LSPs | 121](#)
- [Example: Configuring an LSP Policer | 123](#)
- [Configuring Automatic Policers | 124](#)
- [Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets | 127](#)

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can also configure policers for MPLS LSPs.

The following sections discuss MPLS firewall filters and policers:

Configuring MPLS Firewall Filters

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can then apply this filter to a specific interface. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached. You cannot apply MPLS firewall filters to Ethernet (fxp0) or loopback (lo0) interfaces.

You can configure the following match criteria attributes for MPLS filters at the **[edit firewall family mpls filter *filter-name* term *term-name* from]** hierarchy level:

- **exp**
- **exp-except**

These attributes can accept EXP bits in the range 0 through 7. You can configure the following choices:

- A single EXP bit—for example, **exp 3**;
- Several EXP bits—for example, **exp 0, 4**;
- A range of EXP bits—for example, **exp [0-5]**;

If you do not specify a match criterion (that is, you do not configure the **from** statement and use only the **then** statement with the **count** action keyword), all the MPLS packets passing through the interface on which the filter is applied will be counted.

You also can configure any of the following action keywords at the `[edit firewall family mpls filter filter-name term term-name then]` hierarchy level:

- `count`
- `accept`
- `discard`
- `next`
- `policer`

For more information about how to configure firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*. For more information about how to configure interfaces, see the *Junos OS Network Interfaces Library for Routing Devices* and the *Junos OS Services Interfaces Library for Routing Devices*.

Examples: Configuring MPLS Firewall Filters

The following examples illustrate how you might configure an MPLS firewall filter and then apply the filter to an interface. This filter is configured to count MPLS packets with EXP bits set to either 0 or 4.

The following shows a configuration for an MPLS firewall filter:

```
[edit firewall]
family mpls {
  filter expf {
    term expt0 {
      from {
        exp 0,4;
      }
      then {
        count counter0;
        accept;
      }
    }
  }
}
```

The following shows how to apply the MPLS firewall filter to an interface:

```
[edit interfaces]
so-0/0/0 {
  mtu 4474;
  encapsulation ppp;
  sonet-options {
```



```

        fcs 32;
    }
    unit 0 {
        point-to-point;
        family mpls {
            filter {
                input expf;
                output expf;
            }
        }
    }
}

```

The MPLS firewall filter is applied to the input and output of an interface (see the **input** and **output** statements in the preceding example).

Configuring Policers for LSPs

MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

NOTE: The PTX10003 router only supports regular LSPs.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

You can configure only those match conditions that apply across all types of traffic. The following are the supported match conditions for LSP policers:

- **forwarding-class**
- **packet-length**

- **interface**
- **interface-set**

To enable a policer on an LSP, first you need to configure a policing filter and then include it in the LSP configuration. For information about how to configure policers, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

To configure a policer for an LSP, specify a filter by including the **filter** option to the **policing** statement:

```
policing {
    filter filter-name;
}
```

You can include the **policing** statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit protocols mpls **static-label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **static-label-switched-path** *lsp-name*]

LSP Policer Limitations

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- LSP policers are not supported on aggregated interfaces.
- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.
- LSP policers work on all T Series routers and on M Series routers that have the Internet Processor II application-specific integrated circuit (ASIC).

NOTE: Starting with Junos OS Release 12.2R2, on T Series routers only, you can configure an LSP policer for a specific LSP to be shared across different protocol family types. To do so, you must configure the *logical-interface-policer* statement at the [edit firewall policer *policer-name*] hierarchy level.

Example: Configuring an LSP Policer

The following example shows how you can configure a policing filter for an LSP:

```
[edit firewall]
policer police-ct1 {
  if-exceeding {
    bandwidth-limit 50m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
policer police-ct0 {
  if-exceeding {
    bandwidth-limit 200m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
family any {
  filter bar {
    term discard-ct0 {
      then {
        policer police-ct0;
        accept;
      }
    }
  }
  term discard-ct1 {
    then {
      policer police-ct1;
      accept;
    }
  }
}
```


Configuring Automatic Policers

IN THIS SECTION

- [Configuring Automatic Policers for LSPs | 125](#)
- [Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs | 125](#)
- [Configuring Automatic Policers for Point-to-Multipoint LSPs | 126](#)
- [Disabling Automatic Policing on an LSP | 126](#)
- [Example: Configuring Automatic Policing for an LSP | 127](#)

Automatic policing of LSPs allows you to provide strict service guarantees for network traffic. Such guarantees are especially useful in the context of Differentiated Services for traffic engineered LSPs, providing better emulation for ATM wires over an MPLS network. For more information about Differentiated Services for LSPs, see [“DiffServ-Aware Traffic Engineering Introduction” on page 1194](#).

Differentiated Services for traffic engineered LSPs allow you to provide differential treatment to MPLS traffic based on the EXP bits. To ensure these traffic guarantees, it is insufficient to simply mark the traffic appropriately. If traffic follows a congested path, the requirements might not be met.

LSPs are guaranteed to be established along paths where enough resources are available to meet the requirements. However, even if the LSPs are established along such paths and are marked properly, these requirements cannot be guaranteed unless you ensure that no more traffic is sent to an LSP than there is bandwidth available.

It is possible to police LSP traffic by manually configuring an appropriate filter and applying it to the LSP in the configuration. However, for large deployments it is cumbersome to configure thousands of different filters. Configuration groups cannot solve this problem either, since different LSPs might have different bandwidth requirements, requiring different filters. To police traffic for numerous LSPs, it is best to configure automatic policers.

When you configure automatic policers for LSPs, a policer is applied to all of the LSPs configured on the router. However, you can disable automatic policing on specific LSPs.

NOTE: When you configure automatic policers for DiffServ-aware traffic engineering LSP, GRES is not supported.

NOTE: You cannot configure automatic policing for LSPs carrying CCC traffic.

The following sections describe how to configure automatic policers for LSPs:

Configuring Automatic Policers for LSPs

To configure automatic policers for standard LSPs (neither DiffServ-aware traffic engineered LSPs nor multiclass LSPs), include the **auto-policing** statement with either the **class all *policer-action*** option or the **class ct0 *policer-action*** option:

```
auto-policing {
  class all policer-action;
  class ct0 policer-action;
}
```

You can include this statement at the following hierarchy levels:

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

You can configure the following policer actions for automatic policers:

- **drop**—Drop all packets.
- **loss-priority-high**—Set the packet loss priority (PLP) to high.
- **loss-priority-low**—Set the PLP to low.

These policer actions are applicable to all types of LSPs. The default policer action is to do nothing.

Automatic policers for LSPs police traffic based on the amount of bandwidth configured for the LSPs. You configure the bandwidth for an LSP using the **bandwidth** statement at the [edit protocols mpls **label-switched-path *lsp-path-name***] hierarchy level. If you have enabled automatic policers on a router, change the bandwidth configured for an LSP, and commit the revised configuration, the change does not take effect on the active LSPs. To force the LSPs to use the new bandwidth allocation, issue a **clear mpls lsp** command.

NOTE: You cannot configure automatic policers for LSPs that traverse aggregated interfaces or Multilink Point-to-Point Protocol (MLPPP) interfaces.

Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs

To configure automatic policers for DiffServ-aware traffic engineering LSPs and for multiclass LSPs, include the **auto-policing** statement:


```

auto-policing {
  class all policer-action;
  class ctnumber policer-action;
}

```

You can include this statement at the following hierarchy levels:

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

You include either the **class all *policer-action*** statement or a **class ctnumber *policer-action*** statement for each of one or more classes (you can configure a different policer action for each class). For a list of the actions that you can substitute for the ***policer-action*** variable, see [“Configuring Automatic Policers for LSPs” on page 125](#). The default policer action is to do nothing.

NOTE: You cannot configure automatic policers for LSPs that traverse aggregated interfaces or MLPPP interfaces.

Configuring Automatic Policers for Point-to-Multipoint LSPs

You can configure automatic policers for point-to-multipoint LSPs by including the **auto-policing** statement with either the **class all *policer-action*** option or the **class ct0 *policer-action*** option. You only need to configure the **auto-policing** statement on the primary point-to-multipoint LSP (for more information on primary point-to-multipoint LSPs, see [“Configuring the Primary Point-to-Multipoint LSP” on page 721](#)). No additional configuration is required on the subLSPs for the point-to-multipoint LSP. Point-to-multipoint automatic policing is applied to all branches of the point-to-multipoint LSP. In addition, automatic policing is applied to any local VRF interfaces that have the same forwarding entry as a point-to-multipoint branch. Feature parity for automatic policers for MPLS point-to-multipoint LSPs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

The automatic policer configuration for point-to-multipoint LSPs is identical to the automatic policer configuration for standard LSPs. For more information, see [“Configuring Automatic Policers for LSPs” on page 125](#).

Disabling Automatic Policing on an LSP

When you enable automatic policing, all of the LSPs on the router or logical system are affected. To disable automatic policing on a specific LSP on a router where you have enabled automatic policing, include the **policing** statement with the **no-auto-policing** option:

```

policing no-auto-policing;

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]

Example: Configuring Automatic Policing for an LSP

Configure automatic policing for a multiclass LSP, specifying different actions for class types **ct0**, **ct1**, **ct2**, and **ct3**.

```
[edit protocols mpls]
diffserv-te {
  bandwidth-model extended-mam;
}
auto-policing {
  class ct1 loss-priority-low;
  class ct0 loss-priority-high;
  class ct2 drop;
  class ct3 loss-priority-low;
}
traffic-engineering bgp-igp;
label-switched-path sample-lsp {
  to 3.3.3.3;
  bandwidth {
    ct0 11;
    ct1 1;
    ct2 1;
    ct3 1;
  }
}
interface fxp0.0 {
  disable;
}
interface t1-0/5/3.0;
interface t1-0/5/4.0;
```

Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets

You can selectively set the DiffServ code point (DSCP) field of MPLS-tagged IPv4 and IPv6 packets to 0 without affecting output queue assignment, and continue to set the MPLS EXP field according to the configured rewrite table, which is based on forwarding classes. You can accomplish this by configuring a firewall filter for the MPLS-tagged packets.

For instructions on how to write different DSCP and EXP values in MPLS-tagged IP packets, see the *Class of Service User Guide (Routers and EX9200 Switches)*. For instructions on how to configure firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

RELATED DOCUMENTATION

| [Firewall Filter Match Conditions for MPLS Traffic](#)

System Log Messages and SNMP Traps for MPLS

Whenever an LSP makes a transition from up to down, or down to up, and whenever an LSP switches from one active path to another, the ingress router generates a system log message and sends an SNMP trap. The following shows a sample system log message:

```
RPD_MPLS_LSP_UP: MPLS LSP sheep1 up on primary(any) Route 192.168.1.1 192.168.1.2
192.168.1.3
RPD_MPLS_LSP_CHANGE: MPLS LSP sheep1 change on primary(any) Route 192.168.1.1
192.168.1.2 192.168.1.3
RPD_MPLS_LSP_DOWN: MPLS LSP sheep1 down on primary(any)
```

For information about the MPLS SNMP traps and the proprietary MPLS MIBs, see the *Network Management and Monitoring Guide*.

System log messages for LSPs are generated by default. To disable the default logging of messages for LSPs, configure the **no-syslog** option under the **log-updown** statement:

```
log-updown {
  no-syslog;
}
```

To generate SNMP traps for LSPs, include the **trap** option to the **log-updown** statement:

```
log-updown {
  trap;
}
```

To generate SNMP traps whenever an LSP path goes down, include the **trap-path-down** option to the **log-updown** statement:

```
log-updown {
  trap-path-down;
}
```


To generate SNMP traps whenever an LSP path comes up, include the **trap-path-up** option to the **log-updown** statement:

```
log-updown {
    trap-path-up;
}
```

To disable the generation of system log messages, include the **no-syslog** option to the **log-updown** statement:

```
log-updown {
    no-syslog;
}
```

To disable the generation of SNMP traps, include the **no-trap** statement:

```
no-trap {
    mpls-lsp-traps;
    rfc3812-traps;
}
```

You can include this statement at the following hierarchy levels:

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

For scalability reasons, only the ingress router generates SNMP traps. By default, MPLS issues traps for all configured LSPs. If you have many LSPs, the number of traps can become quite large. To disable the generation of SNMP traps, configure the **no-trap** statement.

The **no-trap** statement also includes the following options which allow you to block certain categories of MPLS SNMP traps:

- **mpls-lsp-traps**—Blocks the MPLS LSP traps defined in the **jnx-mpls.mib**, but allows the **rfc3812.mib** traps.
- **rfc-3812-traps**—Blocks the traps defined in the **rfc3812.mib**, but allows the MPLS LSP traps defined in the **jnx-mpls.mib**.

Configuring Graceful Restart for MPLS

IN THIS CHAPTER

- [Configuring MPLS-Signaled LSPs to Use GRE Tunnels | 131](#)
- [Graceful Restart and MPLS-Related Protocols | 133](#)

Configuring MPLS-Signaled LSPs to Use GRE Tunnels

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

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

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

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

Example: Configuring MPLS-Signaled LSPs to Use GRE Tunnels

To configure MPLS over GRE tunnels:

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

```
[edit interfaces]
interface gr-1/2/0 {
```



```

unit 0 {
    tunnel {
        source 192.168.1.1;
        destination 192.168.1.2;
    }
    family inet {
        address 5.1.1.1/30;
    }
    family iso;
    family mpls;
}

```

2. Enable RSVP and MPLS over the GRE tunnel:

```

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

```

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

```

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

```

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

Graceful Restart and MPLS-Related Protocols

IN THIS SECTION

- [LDP | 133](#)
- [RSVP | 134](#)
- [CCC and TCC | 134](#)

This section contains the following topics:

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*](#)

Configuring IP-Based Filtering for MPLS

IN THIS CHAPTER

- [Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic | 135](#)

Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic

IN THIS SECTION

- [IP-Based Filtering of MPLS Traffic | 135](#)
- [Selective Port Mirroring of MPLS Traffic | 137](#)
- [Sample Configurations | 138](#)

In an MPLS packet, the IP header comes immediately after the MPLS header. The IP-based filtering feature provides a deep inspection mechanism, where a maximum of upto eight MPLS labels of the inner payload can be inspected to enable filtering of MPLS traffic based on IP parameters. The filtered MPLS traffic can also be port mirrored to a monitoring device to offer network-based services in the core MPLS network.

IP-Based Filtering of MPLS Traffic

Prior to Junos OS Release 18.4R1, filtering based on IP parameters was not supported for MPLS family filter. With the introduction of the IP-based filtering feature, you can apply inbound and outbound filters for MPLS-tagged IPv4 and IPv6 packets based on IP parameters, such as source and destination addresses, Layer 4 protocol type, and source and destination ports.

The IP-based filtering feature enables you to filter MPLS packets at the ingress of an interface, where the filtering is done using match conditions on the inner payload of the MPLS packet. The selective MPLS traffic can then be port mirrored to a remote monitoring device using logical tunnels.

To support IP-based filtering, additional match conditions are added that allow MPLS packets to be deep inspected to parse the inner payload with Layer 3 and Layer 4 headers before the appropriate filters are applied.

NOTE: The IP-based filtering feature is supported only for MPLS-tagged IPv4 and IPv6 packets. In other words, the MPLS filters match IP parameters only when the IP payload comes immediately after the MPLS labels.

In other scenarios, where the MPLS payload includes pseudowires, protocols other than inet and inet6, or other encapsulations like Layer 2 VPN or VPLS, the IP-based filtering feature is not supported.

The following match conditions are added for the IP-based filtering of MPLS traffic:

- IPv4 source address
- IPv4 destination address
- IPv6 source address
- IPv6 destination address
- Protocol
- Source port
- Destination port
- Source IPv4 prefix list
- Destination IPv4 prefix list
- Source IPv6 prefix list
- Destination IPv6 prefix list

NOTE: The following match combinations are supported for the IP-based filtering of MPLS traffic:

- Source and destination address match conditions with IPv4 and IPv6 prefix lists.
- Source and destination port address and protocol types match conditions with IPv4 and IPv6 prefix lists.

Selective Port Mirroring of MPLS Traffic

Port mirroring is the capability of mirroring a packet to a configured destination, in addition to the normal processing and forwarding of the packets. Port mirroring is applied as an action for a firewall filter, which is applied at the ingress or egress of any interface. Similarly, the selective port mirroring feature provides the capability to mirror MPLS traffic, which is filtered based on IP parameters, to a mirrored destination using logical tunnels.

To enable selective port mirroring, additional actions are configured at the **[edit firewall family mpls filter filter-name term term-name then]** hierarchy level, in addition to the existing **counter**, **accept**, and **discard** actions:

- **port-mirror**
- **port-mirror-instance**

Port Mirroring

The **port-mirror** action enables port mirroring globally on the device, which applies to all Packet Forwarding Engines (PFEs) and associated interfaces.

For MPLS family filter, the **port-mirror** action is enabled for global port mirroring.

Port Mirroring Instance

The **port-mirror-instance** action enables you to customize each instance with different properties for input sampling and port mirroring output destinations, instead of having to use a single system-wide configuration for port mirroring.

You can configure only two port mirroring instances per Flexible PIC Concentrator (FPC) by including the **instance port-mirror-instance-name** statement at the **[edit forwarding-options port-mirror]** hierarchy level. You can then associate individual port mirroring instances with an FPC, PIC, or (Forwarding Engine Board (FEB) depending on the device hardware.

For MPLS family filter, the **port-mirror-instance** action is enabled only for the port-mirroring instance.

NOTE: For both **port-mirror** and **port-mirror-instance** actions, the output interface must be enabled with Layer 2 family and not family MPLS (Layer 3) for the selective port mirroring feature to work.

Sample Configurations

IN THIS SECTION

- [IP-Based Filtering Configuration | 138](#)
- [Selective Port Mirroring Configuration | 140](#)
- [Mirrored Destination Configuration | 141](#)

IP-Based Filtering Configuration

```
[edit firewall family mpls filter mpls-filter]
term ipv4-term {
  from {
    ip-version {
      ipv4 {
        source-address {
          10.10.10.10/24;
        }
        destination-address {
          20.20.20.20/24;
        }
        protocol tcp {
          source-port 100;
          destination-port 200;
        }
        source-prefix-list ipv4-source-users;
        destination-prefix-list ipv4-destination-users;
      }
    }
    exp 1;
  }
  then port-mirror;
  then accept;
  then count;
}
term ipv6-term {
  from {
    ip-version {
      ipv6 {
        source-address {
          2000::1/128;
        }
      }
    }
  }
}
```



```

    }
    destination-address {
        3000::1/128;
    }
    protocol tcp {
        source-port 100;
        destination-port 200;
    }
    source-prefix-list ipv6-source-users;
    destination-prefix-list ipv6-destination-users;
}
}
exp 1;
}
then port-mirror-instance port-mirror-instance1;
then accept;
then count;
}

```

```

[edit policy-options]
prefix-list ipv4-source-users {
    172.16.1.16/28;
    172.16.2.16/28;
}
prefix-list ipv6-source-users {
    2001::1/128;
    3001::1/128;
}

```

```

[edit interfaces]
xe-0/0/1 {
    unit 0 {
        family inet {
            address 100.100.100.1/30;
        }
        family mpls {
            filter {
                input mpls-filter;
            }
        }
    }
}
}

```


Selective Port Mirroring Configuration

```
[edit forwarding-options]
port-mirroring {
  input {
    rate 2;
    run-length 4;
    maximum-packet-length 500;
  }
  family any {
    output {
      interface xe-2/0/2.0;
    }
  }
}
```

```
[edit forwarding-options]
port-mirroring {
  instance {
    port-mirror-instance1 {
      input {
        rate 3;
        run-length 5;
        maximum-packet-length 500;
      }
      family any {
        output {
          interface xe-2/0/2.0;
        }
      }
    }
  }
}
```

NOTE: The output interface **xe-2/0/2.0** is configured for Layer 2 family and not family MPLS.

For both **port-mirror** and **port-mirror-instance** actions, the output interface must be enabled with Layer 2 family and not family MPLS (Layer 3) for the selective port mirroring feature to work.

Mirrored Destination Configuration

```
[edit interfaces]
xe-2/0/2 {
  vlan-tagging;
  encapsulation extended-vlan-bridge;
  unit 0 {
    vlan-id 600;
  }
}
```

```
[edit bridge-domains]
bd {
  domain-type bridge;
  interface xe-2/0/2.0;
}
```

RELATED DOCUMENTATION

[ip-version](#) | [2637](#)

[ipv4](#) | [2634](#)

[ipv6](#) | [2636](#)

Configuring Link, Node, and Path Protection for MPLS

IN THIS CHAPTER

- Node-Link Protection Overview | 143
- Path Protection Overview | 145
- Configuring Path Protection in an MPLS Network (CLI Procedure) | 145
- Preventing Use of a Path That Previously Failed | 150
- Configuring MPLS Inter-AS Link-Node Protection with Labeled BGP | 151
- Configuring Egress Protection Service Mirroring for BGP Signaled Layer 2 Services | 171
- Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services | 176
- Example: Configuring Layer 3 VPN Egress Protection with PLR as Protector | 198
- Verifying Path Protection in an MPLS Network | 234

Node-Link Protection Overview

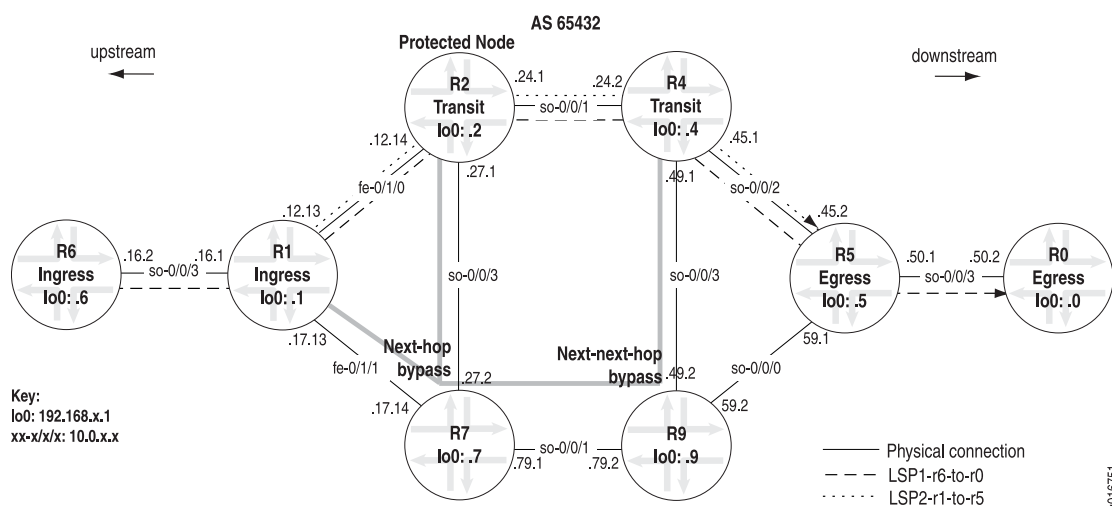
Node-link protection (many-to-one or facility backup) extends the capabilities of link protection and provides slightly different protection from fast reroute. While link protection is useful for selecting an alternate path to the same router when a specific link fails, and fast reroute protects interfaces or nodes along the entire path of an LSP, node-link protection establishes a bypass path that avoids a particular node in the LSP path.

When you enable node-link protection for an LSP, you must also enable link protection on all RSVP interfaces in the path. Once enabled, the following types of bypass paths are established:

- Next-hop bypass LSP—Provides an alternate route for an LSP to reach a neighboring router. This type of bypass path is established when you enable either node-link protection or link protection.
- Next-next-hop bypass LSP—Provides an alternate route for an LSP through a neighboring router en route to the destination router. This type of bypass path is established exclusively when node-link protection is configured.

[Figure 7 on page 144](#) illustrates the example MPLS network topology used in this topic. The example network uses OSPF as the interior gateway protocol (IGP) and a policy to create traffic.

Figure 7: Node-Link Protection



The MPLS network in [Figure 7 on page 144](#) illustrates a router-only network that consists of unidirectional LSPs between **R1** and **R5** (**lsp2-r1-to-r5**) and between **R6** and **R0** (**lsp1-r6-to-r0**). Both LSPs have strict paths configured that go through interface **fe-0/1/0**.

In the network shown in [Figure 7 on page 144](#), both types of bypass paths are preestablished around the protected node (**R2**). A next-hop bypass path avoids interface **fe-0/1/0** by going through **R7**, and a next-next-hop bypass path avoids **R2** altogether by going through **R7** and **R9** to **R4**. Both bypass paths are shared by all protected LSPs traversing the failed link or node (many LSPs protected by one bypass path).

Node-link protection (many-to-one or facility backup) allows a router immediately upstream from a node failure to use an alternate node to forward traffic to its downstream neighbor. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link.

When an outage occurs, the router immediately upstream from the outage switches protected traffic to the bypass node, and then signals the failure to the ingress router. Like fast reroute, node-link protection provides local repair, restoring connectivity faster than the ingress router can establish a standby secondary path or signal a new primary LSP.

Node-link protection is appropriate in the following situations:


- Protection of the downstream link and node is required.
- The number of LSPs to be protected is large.
- Satisfying path selection criteria (priority, bandwidth, and link coloring) for bypass paths is less critical.
- Control at the granularity of individual LSPs is not required.

Path Protection Overview

The main advantages of path protection are control over where the traffic goes after a failure and minimum packet loss when combined with fast reroute (one-to-one backup or link protection). Path protection is the configuration, within a label-switched path (LSP), of two types of paths: a primary path, used in normal operations, and a secondary path used when the primary fails, as shown in [Figure 8 on page 145](#).

In [Figure 8 on page 145](#), an MPLS network consisting of eight routers has a primary path between **R1** and **R5** which is protected by the secondary path between **R1** and **R5**. When a failure is detected, such as an interface down event, an Resource Reservation Protocol (RSVP) error message is sent to the ingress router which switches traffic to the secondary path, maintaining traffic flow.

Figure 8: Path Protection



If the secondary path is pre-sigaled or on standby, recovery time from a failure is faster than if the secondary path is not pre-sigaled. When the secondary path is not pre-sigaled a call-setup delay occurs during which the new physical path for the LSP is established, extending the recovery time. If the failure in the primary path is corrected, and after a few minutes of hold time, the ingress router switches traffic back from the secondary path to the primary path.

Because path protection is provided by the ingress router for the entire path, there can be some disadvantages, for example, double-booking of resources and unnecessary protection of links. By protecting a single resource at a time, local protection can remedy these disadvantages.

Configuring Path Protection in an MPLS Network (CLI Procedure)

The Junos OS implementation of MPLS on EX Series switches provides path protection as a mechanism for protecting against label switched path (LSP) failures. Path protection reduces the time required to recalculate a route in case of a failure within the MPLS tunnel. You configure path protection on the ingress provider edge switch in your MPLS network. You do not configure the egress provider edge switch or the

provider switches for path protection. You can explicitly specify which provider switches are used for the primary and secondary paths, or you can let the software calculate the paths automatically.

Before you configure path protection, be sure you have:

- Configured an ingress provider edge switch and an egress provider edge switch. See [“Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\)” on page 89](#) or [“Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 95](#).
- Configured at least one provider (transit) switch. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).
- Verified the configuration of your MPLS network.

To configure path protection, complete the following tasks on the ingress provider edge switch:

1. [Configuring the Primary Path | 147](#)
2. [Configuring the Secondary Path | 149](#)
3. [Configuring the Revert Timer | 149](#)

Configuring the Primary Path

The **primary** statement creates the primary path, which is the LSP's preferred path. The **secondary** statement creates an alternative path if the primary path can no longer reach the egress provider edge switch.

In the tasks described in this topic, the **lsp-name** has already been configured on the ingress provider edge switch as **lsp_to_240** and the loopback interface address on the remote provider edge switch has already been configured as **127.0.0.8**.

When the software switches from the primary to the secondary path, it continuously attempts to revert to the primary path, switching back to it when it is again reachable but no sooner than the retry time specified in the **revert-timer** statement.

You can configure zero primary paths or one primary path. If you do not configure a primary path, the first secondary path (if a secondary path has been configured) is selected as the path. If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary for the packets to reach the egress provider edge switch.

To configure a primary path:

1. Create the primary path for the LSP:

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set primary primary_path_lsp_to_240
```

2. Configure an explicit route for the primary path by specifying the IP address of the loopback interface or the switch IP address or hostname of each switch used in the MPLS tunnel. You can specify the link types as either **strict** or **loose** in each **path** statement. If the link type is **strict**, the LSP must go to the next address specified in the **path** statement without traversing other switches. If the link type is **loose**, the LSP can traverse through other switches before reaching this switch. This configuration uses the default **strict** designation for the paths.

NOTE: You can enable path protection without specifying which provider switches are used. If you do not list the specific provider switches to be used for the MPLS tunnel, the switch calculates the route.

TIP: Do not include the ingress provider edge switch in these statements. List the IP address of the loopback interface or switch address or hostname of all other switch hops in sequence, ending with the egress provider edge switch.

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set path primary_path_lsp_to_240 127.0.0.2
user@switch# set path primary_path_lsp_to_240 127.0.0.3
```



```
user@switch# set path primary_path_lsp_to_240 127.0.0.8
```

Configuring the Secondary Path

You can configure zero or more secondary paths. All secondary paths are equal, and the software tries them in the order that they are listed in the configuration. The software does not attempt to switch among secondary paths. If the first secondary path in the configuration is not available, the next one is tried, as so on. To create a set of equal paths, specify secondary paths without specifying a primary path. If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary to reach the egress provider edge switch.

To configure the secondary path:

1. Create a secondary path for the LSP:

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set secondary secondary_path_lsp_to_240 standby
```

2. Configure an explicit route for the secondary path by specifying the IP address of the loopback interface or the switch IP address or hostname of each switch used in the MPLS tunnel. You can specify the link types as either **strict** or **loose** in each **path** statement. This configuration uses the default **strict** designation for the paths.

TIP: Do not include the ingress provider edge switch in these statements. List the IP address of the loopback interface or switch address or hostname of all other switch hops in sequence, ending with the egress provider edge switch.

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set path secondary_path_lsp_to_240 127.0.0.4
user@switch# set path primary_path_lsp_to_240 127.0.0.8
```

Configuring the Revert Timer

For LSPs configured with both primary and secondary paths, you can optionally configure a revert timer. If the primary path goes down and traffic is switched to the secondary path, the revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert traffic back to the primary path. If the primary path experiences any connectivity problems or stability problems during this time, the timer is restarted.

TIP: If you do not explicitly configure the revert timer, it is set by default to 60 seconds.

To configure the revert timer for LSPs configured with primary and secondary paths:

- For all LSPs on the switch:

```
[edit protocols mpls]
user@switch# set revert-timer 120
```

- For a specific LSP on the switch:

```
[edit protocols mpls label-switched-path]
user@switch# set lsp_to_240 revert-timer 120
```

RELATED DOCUMENTATION

[Understanding MPLS and Path Protection on EX Series Switches](#) | 43

Preventing Use of a Path That Previously Failed

If you configure an alternate path through the network in case the active path fails, you may not want traffic to revert back to the failed path, even if it is no longer failing. When you configure a primary path, the traffic switches over to the secondary path during a failure, and reverts back to the primary path when it returns.

At times, switching traffic back to a primary path that has previously failed may not be a particularly sound idea. In this case, only configure secondary paths, resulting in the next configured secondary path establishing when the first secondary path fails. Later, if the first secondary path becomes operational, the Junos OS will not revert to it, but will continue using the second secondary path.

Configuring MPLS Inter-AS Link-Node Protection with Labeled BGP

IN THIS SECTION

- [Understanding MPLS Inter-AS Link Protection | 151](#)
- [Example: Configuring MPLS Inter-AS Link-Node Protection | 153](#)

Understanding MPLS Inter-AS Link Protection


Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. The ingress router chooses an alternate link through another interface to send traffic to its destination.

In [Figure 9 on page 152](#), autonomous system border routers (ASBRs) run external BGP (EBGP) to ASBRs in another autonomous system (AS) to exchange labels for /32 IPv4 routes. Inside the ASs, internal BGP (IBGP) propagates the routes to provider edge (PE) devices. If the link from Device ASBR3 to Device ASBR1 goes down, until Device ASBR3 reinstalls the new next hop, all traffic going toward AS 64510 from AS 64511 through the ASBR3-ASBR1 link is dropped. A fast traffic restoration can be achieved if Device ASBR3 preprograms a backup path either through Device ASBR4 or through a direct path to Device ASBR2 if one exists (not shown in the diagram). This assumes that Device ASBR3 learns a loop-free MPLS path for routes that need to be protected either through IBGP or EBGP.

This solution does not handle a failure on Device ASBR3 for traffic going toward AS 64511 from AS 64510 through the ASBR3-ASBR1 link. This solution is limited to downstream inter-AS link-node protection with labeled BGP. This solution does not support service restoration between provider (P) and ASBR routers when there is an ASBR failure. For example, this solution does not handle a failure on the P3-ASBR3 link.

This supported functionality is similar to BGP multipath, except only one next hop is used for active forwarding, and a second path is in protected mode.

Figure 9: MPLS Inter-AS Link-Node Protection Conceptual Topology



In an MPLS inter-AS environment, link protection can be enabled when **labeled-unicast** is used to send traffic between ASs. Hence, MPLS inter-AS link protection is configured on the link between two routers in different ASs.

To configure link protection on an interface, use the **protection** statement at the **[edit protocols bgp group group-name family inet labeled-unicast]** hierarchy level:

```
protocols {
  bgp {
    group test1 {
      type external;
      local-address 192.168.1.2;
      family inet {
        labeled-unicast {
          protection;
        }
      }
    }
  }
}
```


NOTE: MPLS inter-AS link protection is supported only with **labeled-unicast** and external peers in a master routing instance.

The link on which protection is configured is known as the protection path. A protection path is selected only after the best path selection and is not selected in the following cases:

- The best path is a non-BGP path.
- Multiple next hops are active, as in BGP multipath.

SEE ALSO

[Example: Configuring MPLS Inter-AS Link-Node Protection | 153](#)

Example: Configuring MPLS Inter-AS Link-Node Protection

IN THIS SECTION

- [Requirements | 153](#)
- [Overview | 153](#)
- [Configuration | 154](#)
- [Verification | 167](#)

This example shows how to configure tail-end protection in an inter-AS deployment with Layer 3 VPNs.

Requirements

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

Overview

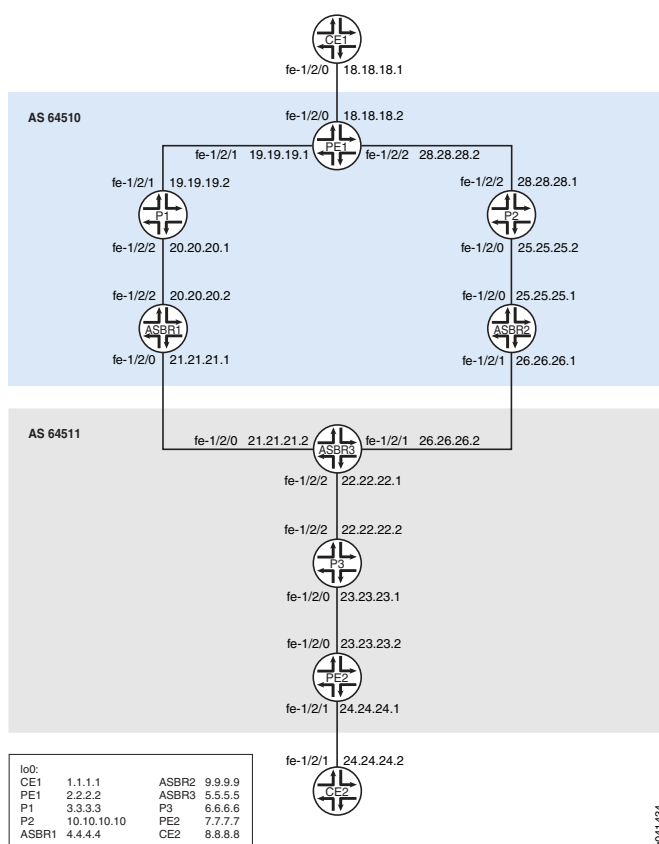
In [Figure 10 on page 154](#), autonomous system border routers (ASBRs) run external BGP (EBGP) to ASBRs in another autonomous system (AS) to exchange labels for /32 IPv4 routes. Inside the ASs, internal BGP (IBGP) propagates the routes to provider edge (PE) devices.

If the link from Device ASBR3 to Device ASBR1 goes down, until ASBR3 reinstalls the new next hop, all traffic going toward AS 64510 from AS 64511 through the ASBR3-ASBR1 link is dropped.

This example shows how to achieve fast traffic restoration by configuring Device ASBR3 to preprogram a backup path through Device ASBR2.

NOTE: This solution does not handle the Device P3 to Device ASBR3 failure. Nor does it handle a failure on Device ASBR3 for traffic going toward AS 64511 from AS 64510 through the ASBR3-ASBR1 link. This traffic is dropped.

Figure 10: MPLS Inter-AS Link-Node Protection Example Topology



Configuration

CLI Quick Configuration

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

Device ASBR1


```

set interfaces fe-1/2/2 unit 0 family inet address 20.20.20.2/30
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/2/0 unit 0 family inet address 21.21.21.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 4.4.4.4/32
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path To_PE1 to 2.2.2.2
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols bgp group To-PE1 type internal
set protocols bgp group To-PE1 local-address 4.4.4.4
set protocols bgp group To-PE1 family inet unicast
set protocols bgp group To-PE1 family inet labeled-unicast
set protocols bgp group To-PE1 export next-hop-self
set protocols bgp group To-PE1 neighbor 2.2.2.2 family inet labeled-unicast
set protocols bgp group To-ASBR3 type external
set protocols bgp group To-ASBR3 family inet labeled-unicast
set protocols bgp group To-ASBR3 export To-ASBR3
set protocols bgp group To-ASBR3 neighbor 21.21.21.2 peer-as 64511
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options policy-statement To-ASBR3 term 1 from route-filter 2.2.2.2/32 exact
set policy-options policy-statement To-ASBR3 term 1 then accept
set policy-options policy-statement To-ASBR3 term 2 then reject
set policy-options policy-statement next-hop-self then next-hop self
set routing-options autonomous-system 64510

```

Device ASBR2

```

set interfaces fe-1/2/0 unit 0 description to-P2
set interfaces fe-1/2/0 unit 0 family inet address 25.25.25.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 description to-ASBR3
set interfaces fe-1/2/1 unit 0 family inet address 26.26.26.1/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 9.9.9.9/32

```



```

set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path To_PE1 to 2.2.2.2
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.0
set protocols mpls interface lo0.0
set protocols bgp group To-PE1 type internal
set protocols bgp group To-PE1 local-address 9.9.9.9
set protocols bgp group To-PE1 family inet unicast
set protocols bgp group To-PE1 family inet labeled-unicast
set protocols bgp group To-PE1 export next-hop-self
set protocols bgp group To-PE1 neighbor 2.2.2.2 family inet labeled-unicast
set protocols bgp group To-ASBR3 type external
set protocols bgp group To-ASBR3 family inet labeled-unicast
set protocols bgp group To-ASBR3 export To-ASBR3
set protocols bgp group To-ASBR3 neighbor 26.26.26.2 peer-as 64511
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options policy-statement To-ASBR3 term 1 from route-filter 2.2.2.2/32 exact
set policy-options policy-statement To-ASBR3 term 1 then accept
set policy-options policy-statement To-ASBR3 term 2 then reject
set policy-options policy-statement next-hop-self then next-hop self
set routing-options autonomous-system 64510

```

Device ASBR3

```

set interfaces fe-1/2/0 unit 0 description to-ASBR1
set interfaces fe-1/2/0 unit 0 family inet address 21.21.21.2/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 description to-P3
set interfaces fe-1/2/2 unit 0 family inet address 22.22.22.1/30
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 description to-ASBR2
set interfaces fe-1/2/1 unit 0 family inet address 26.26.26.2/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 5.5.5.5/32
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0

```



```

set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path To_PE2 to 7.7.7.7
set protocols mpls interface lo0.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/1.0
set protocols bgp group To-PE2 type internal
set protocols bgp group To-PE2 local-address 5.5.5.5
set protocols bgp group To-PE2 family inet unicast
set protocols bgp group To-PE2 export next-hop-self
set protocols bgp group To-PE2 neighbor 7.7.7.7 family inet labeled-unicast
set protocols bgp group To-ASBR1 type external
set protocols bgp group To-ASBR1 family inet labeled-unicast protection
set protocols bgp group To-ASBR1 family inet labeled-unicast per-prefix-label
set protocols bgp group To-ASBR1 export To-ASBR1
set protocols bgp group To-ASBR1 neighbor 21.21.21.1 peer-as 64510
set protocols bgp group To-ASBR2 type external
set protocols bgp group To-ASBR2 family inet labeled-unicast protection
set protocols bgp group To-ASBR2 family inet labeled-unicast per-prefix-label
set protocols bgp group To-ASBR2 export To-ASBR2
set protocols bgp group To-ASBR2 neighbor 26.26.26.1 peer-as 64510
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/1.0
set policy-options policy-statement To-ASBR1 term 1 from route-filter 7.7.7.7/32 exact
set policy-options policy-statement To-ASBR1 term 1 then accept
set policy-options policy-statement To-ASBR1 term 2 then reject
set policy-options policy-statement To-ASBR2 term 1 from route-filter 7.7.7.7/32 exact
set policy-options policy-statement To-ASBR2 term 1 then accept
set policy-options policy-statement To-ASBR2 term 2 then reject
set policy-options policy-statement next-hop-self then next-hop self
set routing-options autonomous-system 64511

```

Device CE1

```

set interfaces fe-1/2/0 unit 0 family inet address 18.18.18.1/30
set interfaces lo0 unit 0 family inet address 1.1.1.1/32

```



```
set protocols ospf area 0.0.0.2 interface fe-1/2/0.0
set protocols ospf area 0.0.0.2 interface lo0.0 passive
```

Device CE2

```
set interfaces fe-1/2/1 unit 0 family inet address 24.24.24.2/30
set interfaces lo0 unit 0 family inet address 8.8.8.8/32
set protocols bgp group To_PE2 neighbor 24.24.24.1 export myroutes
set protocols bgp group To_PE2 neighbor 24.24.24.1 peer-as 64511
set policy-options policy-statement myroutes from protocol direct
set policy-options policy-statement myroutes then accept
set routing-options autonomous-system 64509
```

Device P1

```
set interfaces fe-1/2/1 unit 0 family inet address 19.19.19.2/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 family inet address 20.20.20.1/30
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 3.3.3.3/32
set protocols rsvp interface fe-1/2/1.0
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols mpls interface fe-1/2/1.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/1.0
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Device P2

```
set interfaces fe-1/2/0 unit 0 description to-ASBR2
```



```

set interfaces fe-1/2/0 unit 0 family inet address 25.25.25.2/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 description to-PE1
set interfaces fe-1/2/2 unit 0 family inet address 28.28.28.1/30
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.10.10.10/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Device P3

```

set interfaces fe-1/2/2 unit 0 family inet address 22.22.22.2/30
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/2/0 unit 0 family inet address 23.23.23.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 6.6.6.6/32
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Device PE1


```

set interfaces fe-1/2/0 unit 0 family inet address 18.18.18.2/30
set interfaces fe-1/2/1 unit 0 family inet address 19.19.19.1/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 description to-P2
set interfaces fe-1/2/2 unit 0 family inet address 28.28.28.2/30
set interfaces lo0 unit 0 family inet address 2.2.2.2/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/2.0
set protocols mpls label-switched-path To-ASBR1 to 4.4.4.4
set protocols mpls label-switched-path To-ASBR2 to 9.9.9.9
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fe-1/2/2.0
set protocols bgp group To_ASBR1 type internal
set protocols bgp group To_ASBR1 local-address 2.2.2.2
set protocols bgp group To_ASBR1 family inet labeled-unicast
set protocols bgp group To_ASBR1 neighbor 4.4.4.4 family inet labeled-unicast resolve-vpn
set protocols bgp group To_PE2 type external
set protocols bgp group To_PE2 multihop ttl 20
set protocols bgp group To_PE2 local-address 2.2.2.2
set protocols bgp group To_PE2 family inet-vpn unicast
set protocols bgp group To_PE2 neighbor 7.7.7.7 peer-as 64511
set protocols bgp group To_ASBR2 type internal
set protocols bgp group To_ASBR2 local-address 2.2.2.2
set protocols bgp group To_ASBR2 family inet labeled-unicast
set protocols bgp group To_ASBR2 neighbor 9.9.9.9 family inet labeled-unicast resolve-vpn
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
set policy-options policy-statement bgp-to-ospf term 2 then reject
set policy-options policy-statement vpnexport term 1 from protocol ospf
set policy-options policy-statement vpnexport term 1 then community add test_comm
set policy-options policy-statement vpnexport term 1 then accept
set policy-options policy-statement vpnexport term 2 then reject
set policy-options policy-statement vpnimport term 1 from protocol bgp
set policy-options policy-statement vpnimport term 1 from community test_comm
set policy-options policy-statement vpnimport term 1 then accept
set policy-options policy-statement vpnimport term 2 then reject

```



```

set policy-options community test_comm members target:1:64510
set routing-instances vpn2CE1 instance-type vrf
set routing-instances vpn2CE1 interface fe-1/2/0.0
set routing-instances vpn2CE1 route-distinguisher 1:64510
set routing-instances vpn2CE1 vrf-import vpnimport
set routing-instances vpn2CE1 vrf-export vpnexport
set routing-instances vpn2CE1 protocols ospf export bgp-to-ospf
set routing-instances vpn2CE1 protocols ospf area 0.0.0.2 interface fe-1/2/0.0
set routing-options autonomous-system 64510

```

Device PE2

```

set interfaces fe-1/2/0 unit 0 family inet address 23.23.23.2/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 24.24.24.1/30
set interfaces lo0 unit 0 family inet address 7.7.7.7/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols mpls label-switched-path To-ASBR3 to 5.5.5.5
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols bgp group To_ASBR3 type internal
set protocols bgp group To_ASBR3 local-address 7.7.7.7
set protocols bgp group To_ASBR3 family inet labeled-unicast
set protocols bgp group To_ASBR3 neighbor 5.5.5.5 family inet labeled-unicast resolve-vpn
set protocols bgp group To_PE1 type external
set protocols bgp group To_PE1 multihop ttl 20
set protocols bgp group To_PE1 local-address 7.7.7.7
set protocols bgp group To_PE1 family inet-vpn unicast
set protocols bgp group To_PE1 neighbor 2.2.2.2 peer-as 64510
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options policy-statement vpnexport term 1 from protocol bgp
set policy-options policy-statement vpnexport term 1 then community add test_comm
set policy-options policy-statement vpnexport term 1 then accept
set policy-options policy-statement vpnexport term 2 then reject
set policy-options policy-statement vpnimport term 1 from protocol bgp
set policy-options policy-statement vpnimport term 1 from community test_comm
set policy-options policy-statement vpnimport term 1 then accept

```



```

set policy-options policy-statement vpnimport term 2 then reject
set policy-options community test_comm members target:1:64510
set routing-instances vpn2CE2 instance-type vrf
set routing-instances vpn2CE2 interface fe-1/2/1.0
set routing-instances vpn2CE2 route-distinguisher 1:64510
set routing-instances vpn2CE2 vrf-import vpnimport
set routing-instances vpn2CE2 vrf-export vpnexport
set routing-instances vpn2CE2 protocols bgp group To_CE2 peer-as 64509
set routing-instances vpn2CE2 protocols bgp group To_CE2 neighbor 24.24.24.2
set routing-options autonomous-system 64511

```

Step-by-Step Procedure

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

To configure the EBGp scenario:

1. Configure the router interfaces.

```

[edit interfaces]
user@ASBR3# set fe-1/2/0 unit 0 description to-ASBR1
user@ASBR3# set fe-1/2/0 unit 0 family inet address 21.21.21.2/30
user@ASBR3# set fe-1/2/0 unit 0 family mpls
user@ASBR3# set fe-1/2/2 unit 0 description to-P3
user@ASBR3# set fe-1/2/2 unit 0 family inet address 22.22.22.1/30
user@ASBR3# set fe-1/2/2 unit 0 family mpls
user@ASBR3# set fe-1/2/1 unit 0 description to-ASBR2
user@ASBR3# set fe-1/2/1 unit 0 family inet address 26.26.26.2/30
user@ASBR3# set fe-1/2/1 unit 0 family mpls
user@ASBR3# set lo0 unit 0 family inet address 5.5.5.5/32

```

2. Configure an interior gateway protocol (IGP), such as OSPF or IS-IS.

```

[edit protocols ospf]
user@ASBR3# set traffic-engineering
[edit protocols ospf area 0.0.0.0]
user@ASBR3# set interface fe-1/2/2.0
user@ASBR3# set interface lo0.0 passive
user@ASBR3# set interface fe-1/2/1.0

```


3. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@ASBR3# set autonomous-system 64511
```

4. Configure the routing policy.

```
[edit policy-options policy-statement To-ASBR1]
user@ASBR3# set term 1 from route-filter 7.7.7.7/32 exact
user@ASBR3# set term 1 then accept
user@ASBR3# set term 2 then reject
[edit policy-options policy-statement To-ASBR2]
user@ASBR3# set term 1 from route-filter 7.7.7.7/32 exact
user@ASBR3# set term 1 then accept
user@ASBR3# set term 2 then reject
[edit policy-options policy-statement next-hop-self]
user@ASBR3# set then next-hop self
```

5. Configure the EBGP sessions.

```
[edit protocols bgp group To-ASBR1]
user@ASBR3# set type external
user@ASBR3# set family inet labeled-unicast protection
user@ASBR3# set family inet labeled-unicast per-prefix-label
user@ASBR3# set export To-ASBR1
user@ASBR3# set neighbor 21.21.21.1 peer-as 64510
[edit protocols bgp group To-ASBR2]
user@ASBR3# set type external
user@ASBR3# set family inet labeled-unicast protection
user@ASBR3# set family inet labeled-unicast per-prefix-label
user@ASBR3# set export To-ASBR2
user@ASBR3# set neighbor 26.26.26.1 peer-as 64510
```

6. Configure the IBGP sessions.

```
[edit protocols bgp group To-PE2]
user@ASBR3# set type internal
user@ASBR3# set local-address 5.5.5.5
user@ASBR3# set family inet unicast
user@ASBR3# set export next-hop-self
user@ASBR3# set neighbor 7.7.7.7 family inet labeled-unicast
```


7. Configure MPLS.

```
[edit protocols mpls]
user@ASBR3# set traffic-engineering bgp-igp-both-ribs
user@ASBR3# set label-switched-path To_PE2 to 7.7.7.7
user@ASBR3# set interface lo0.0
user@ASBR3# set interface fe-1/2/0.0
user@ASBR3# set interface fe-1/2/2.0
user@ASBR3# set interface fe-1/2/1.0
```

8. Configure a signaling protocol.

```
[edit protocols rsvp]
user@ASBR3# set interface fe-1/2/2.0
user@ASBR3# set interface lo0.0
user@ASBR3# set interface fe-1/2/0.0
user@ASBR3# set interface fe-1/2/1.0
```

Results

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

```
user@ASBR3# show interfaces
fe-1/2/0 {
  unit 0 {
    description to-ASBR1;
    family inet {
      address 21.21.21.2/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    description to-ASBR2;
    family inet {
      address 26.26.26.2/30;
    }
    family mpls;
  }
}
```



```

fe-1/2/2 {
  unit 0 {
    description to-P3;
    family inet {
      address 22.22.22.1/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 5.5.5.5/32;
    }
  }
}

```

```

user@ASBR3# show protocols
rsvp {
  interface fe-1/2/2.0;
  interface lo0.0;
  interface fe-1/2/0.0;
  interface fe-1/2/1.0;
}
mpls {
  traffic-engineering bgp-igp-both-ribs;
  label-switched-path To_PE2 {
    to 7.7.7.7;
  }
  interface lo0.0;
  interface fe-1/2/0.0;
  interface fe-1/2/2.0;
  interface fe-1/2/1.0;
}
bgp {
  group To-PE2 {
    type internal;
    local-address 5.5.5.5;
    family inet {
      unicast;
    }
    export next-hop-self;
    neighbor 7.7.7.7 {
      family inet {

```



```

        labeled-unicast;
    }
}
}
group To-ASBR1 {
    type external;
    family inet {
        labeled-unicast {
            protection;
        }
    }
    export To-ASBR1;
    neighbor 21.21.21.1 {
        peer-as 64510;
    }
}
group To-ASBR2 {
    type external;
    family inet {
        labeled-unicast {
            protection;
        }
    }
    export To-ASBR2;
    neighbor 26.26.26.1 {
        peer-as 64510;
    }
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-1/2/2.0;
        interface lo0.0 {
            passive;
        }
        interface fe-1/2/1.0;
    }
}
}

```

```

user@ASBR3# show policy-options
policy-statement To-ASBR1 {
    term 1 {
        from {

```



```

        route-filter 7.7.7.7/32 exact;
    }
    then accept;
}
term 2 {
    then reject;
}
}
policy-statement To-ASBR2 {
    term 1 {
        from {
            route-filter 7.7.7.7/32 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement next-hop-self {
    then {
        next-hop self;
    }
}

```

```

user@ASBR3# show routing-options
autonomous-system 64511;

```

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

Verification

IN THIS SECTION

- [Checking the BGP Neighbor Sessions | 167](#)
- [Checking the Routes | 170](#)

Confirm that the configuration is working properly.

Checking the BGP Neighbor Sessions

Purpose

Verify that BGP protection is enabled.

Action

user@ASBR3# **show bgp neighbor 21.21.21.1**

```

Peer: 21.21.21.1+58259 AS 64510 Local: 21.21.21.2+179 AS 64511
  Type: External      State: Established      Flags: <ImportEval Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ To-ASBR1 ]
  Options: <Preference AddressFamily PeerAS Refresh>
  Options: <Protection>
  Address families configured: inet-labeled-unicast
  Holdtime: 90 Preference: 170
NLRI configured with protection: inet-labeled-unicast
  Number of flaps: 0
  Peer ID: 4.4.4.4      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Group index: 4      Peer index: 0
  BFD: disabled, down
  Local Interface: fe-1/2/0.0
  NLRI for restart configured on peer: inet-labeled-unicast
  NLRI advertised by peer: inet-labeled-unicast
  NLRI for this session: inet-labeled-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-labeled-unicast
  NLRI of received end-of-rib markers: inet-labeled-unicast
  NLRI of all end-of-rib markers sent: inet-labeled-unicast
  Peer supports 4 byte AS extension (peer-as 64510)
  Peer does not support Addpath
  Table inet.0 Bit: 10001
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        1
    Accepted prefixes:        1
    Suppressed due to damping: 0
    Advertised prefixes:      1
  Last traffic (seconds): Received 7      Sent 20      Checked 32
  Input messages:  Total 170      Updates 2      Refreshes 0      Octets 3326

```



```
Output messages: Total 167      Updates 1      Refreshes 0      Octets 3288
Output Queue[0]: 0
```

user@ASBR3# **show bgp neighbor 26.26.26.1**

```
Peer: 26.26.26.1+61072 AS 64510 Local: 26.26.26.2+179 AS 64511
  Type: External      State: Established      Flags: <ImportEval Sync>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Export: [ To-ASBR2 ]
  Options: <Preference AddressFamily PeerAS Refresh>
  Options: <Protection>
  Address families configured: inet-labeled-unicast
  Holdtime: 90 Preference: 170
NLRI configured with protection: inet-labeled-unicast
  Number of flaps: 0
  Peer ID: 9.9.9.9      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Group index: 5      Peer index: 0
  BFD: disabled, down
  Local Interface: fe-1/2/1.0
  NLRI for restart configured on peer: inet-labeled-unicast
  NLRI advertised by peer: inet-labeled-unicast
  NLRI for this session: inet-labeled-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-labeled-unicast
  NLRI of received end-of-rib markers: inet-labeled-unicast
  NLRI of all end-of-rib markers sent: inet-labeled-unicast
  Peer supports 4 byte AS extension (peer-as 64510)
  Peer does not support Addpath
  Table inet.0 Bit: 10002
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:      1
    Received prefixes:    1
    Accepted prefixes:    1
    Suppressed due to damping: 0
    Advertised prefixes:  1
  Last traffic (seconds): Received 21      Sent 9      Checked 42
  Input messages: Total 170      Updates 2      Refreshes 0      Octets 3326
  Output messages: Total 168      Updates 1      Refreshes 0      Octets 3307
  Output Queue[0]: 0
```


Meaning

The output shows that the **Protection** option is enabled for the EBGp peers, Device ASBR1 and Device ASBR2.

This is also shown with the **NLRI configured with protection: inet-labeled-unicast** screen output.

Checking the Routes

Purpose

Make sure that the backup path is installed in the routing table.

Action

user@ASBR3> **show route 2.2.2.2**

```
inet.0: 12 destinations, 14 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          *[BGP/170] 01:36:25, MED 2, localpref 100
                   AS path: 64510 I, validation-state: unverified
                   > to 21.21.21.1 via fe-1/2/0.0, Push 299824
                   to 26.26.26.1 via fe-1/2/1.0, Push 299808
                   [BGP/170] 01:36:25, MED 2, localpref 100
                   AS path: 64510 I, validation-state: unverified
                   > to 26.26.26.1 via fe-1/2/1.0, Push 299808
```

Meaning

The **show route** command displays active as well as backup paths to Device PE1.

SEE ALSO

[Understanding MPLS Inter-AS Link Protection | 151](#)

Example: Preventing BGP Session Resets

Examples: Configuring BGP Flap Damping

RELATED DOCUMENTATION

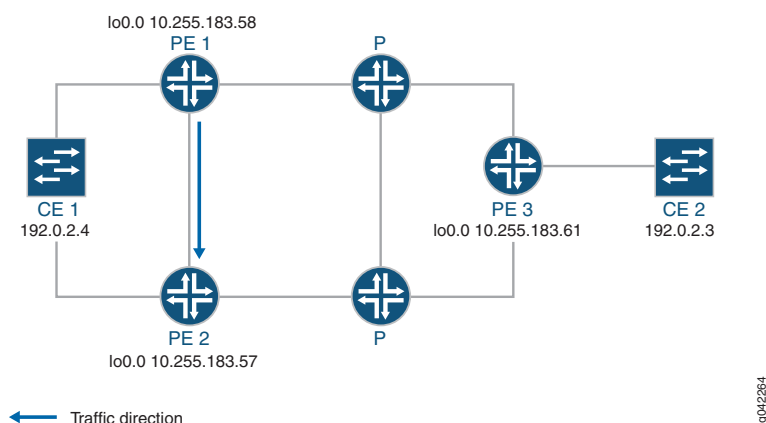
Example: Configuring Provider Edge Link Protection in Layer 3 VPNs

Configuring Egress Protection Service Mirroring for BGP Signaled Layer 2 Services

Starting in Junos OS Release 14.2, Junos OS supports the restoration of egress traffic when there is a link or node failure in the egress PE node. If there is a link or node failure in the core network, a protection mechanism such as MPLS fast reroute can be triggered on the transport LSPs between the PE routers to repair the connection within tens of milliseconds. An egress protection LSP addresses the problem of a node-link failure at the edge of the network (for example, a failure of a PE router).

Figure 1 shows a simplified topology of the use case that explains this feature.

Figure 11: Egress Protection LSP Configured from Router PE1 to Router PE2



CE1 is multihomed to PE1 and PE2. There are two paths connecting CE1 and CE2. The working path is CE2-PE3-P-PE1-CE1, via pseudowire PW21. The protecting path is CE2-PE3-P-PE2-CE1, via pseudowire PW22. Traffic is flowing through the working path under normal circumstances. When the end-to-end OAM between CE1 and CE2 detects failure on the working path, traffic will be switched from the working path to the protecting path. The end-to-end failure detection and recovery relies on control plane hence should be relatively slow. To achieve faster protection, local repair mechanisms similar to those used by MPLS fast reroute should be used. In Figure 1 above, if link or node failed in the core network (like link failure on P-PE1, P-PE3, or node failure on P), the MPLS fast reroute will happen on the transport LSPs between PE1 and PE3. The failure could be locally repaired within tens of milliseconds. However, if link or node failure happens at the edge (like link failure on PE3-CE2 or node failure on PE3), there is no local repair currently so we have to rely on the CE1-CE2 end-to-end protection to repair the failure.

- Device CE2—Traffic origin
- Router PE3—Ingress PE router
- Router PE1— (Primary) Egress PE router

- Router PE2—Protector PE router
- Device CE1—Traffic destination

When the link between CE1– PE1 goes down, PE1 will briefly redirect that traffic towards CE1, to PE2. PE2 forwards it to CE1 until ingress router PE3 recalculates to forward the traffic to PE2.

Initially the traffic direction was; CE2 – PE3 – P – PE1 – CE1.

When the link between CE1– PE1 goes down, the traffic will be; CE2 – PE3 – P – PE1 – PE2 –CE1. PE3 then recalculates the path; CE2 – PE3 – P – PE2 – CE1.

1. Configure RSVP on PE1, PE2, and PE3.

```
[edit protocols]
user@PE1# set interface all
user@PE2# set interface all
user@PE3# set interface all
```

2. Configure MPLS.

```
[edit protocols mpls]
user@PE1# set interface all
user@PE2# set interface all
user@PE3# set interface all
```

3. Set PE1 as **primary** and PE2 as **protector** nodes.

```
[edit protocols mpls]
user@PE1# set egress-protection context-identifier address primary
user@PE2# set egress-protection context-identifier address protector
```

4. Enable **egress-protection** on PE1 and PE2.

```
[edit protocols bgp]
user@PE1# set group ibgp family l2vpn egress-protection
user@PE2# set group ibgp family l2vpn egress-protection
```

5. Configure LDP and ISIS on PE1, PE2, and PE3.

```
[edit protocols ldp]
user@PE1# set interface all
```



```
user@PE2# set interface all
user@PE3# set interface all
```

```
[edit protocols isis]
user@PE1# set interface all point-to-point
user@PE2# set interface all point-to-point
user@PE3# set interface all point-to-point
```

6. Configure a load balancing policy at PE1, PE2, and PE3.

```
[edit]
user@PE1# set policy-options policy-statement lb then load-balance per-packet
user@PE2# set policy-options policy-statement lb then load-balance per-packet
user@PE3# set policy-options policy-statement lb then load-balance per-packet
```

7. Configure the routing options at PE1, PE2, and PE3, to export routes based on the load balancing policy.

```
[edit]
user@PE1# set routing-options traceoptions file ro.log
user@PE1# set routing-options traceoptions flag normal
user@PE1# set routing-options traceoptions flag route
user@PE1# set routing-options autonomous-system 100
user@PE1# set routing-options forwarding-table export lb
```

```
[edit]
user@PE2# set routing-options traceoptions file ro.log
user@PE2# set routing-options traceoptions flag normal
user@PE2# set routing-options traceoptions flag route
user@PE2# set routing-options autonomous-system 100
user@PE2# set routing-options forwarding-table export lb
```

```
[edit]
user@PE3# set routing-options traceoptions file ro.log
user@PE3# set routing-options traceoptions flag normal
user@PE3# set routing-options traceoptions flag route
user@PE3# set routing-options autonomous-system 100
user@PE3# set routing-options forwarding-table export lb
```

8. Configure BGP at PE1 to advertise nrli from the routing instance with context-ID as next-hop.


```
[edit]
user@PE1# set routing-instances foo egress-protection context-identifier context-identifier
```

9. Configure l2vpn at PE1, PE2, and PE3

At PE1:

```
[edit routing-instances]
foo {
  instance-type l2vpn;
  egress-protection {
    context-identifier {
      198.51.100.0;
    }
  }
  interface ge-2/0/2.0;
  route-distinguisher 10.255.183.58:1;
  vrf-target target:9000:1;
  protocols {
    l2vpn {
      encapsulation-type ethernet-vlan;
      site foo {
        site-identifier 1;
        multi-homing;
        site-preference primary;
        interface ge-2/0/2.0 {
          remote-site-id 2;
        }
      }
    }
  }
}
```

At PE2:

```
[edit routing-instances]
foo {
  instance-type l2vpn;
  egress-protection {
    protector;
  }
  interface ge-2/0/2.0;
  route-distinguisher 10.255.183.57:1;
```



```
vrf-target target:9000:1;
protocols {
  l2vpn {
    encapsulation-type ethernet-vlan;
    site foo{
      site-identifier 1;
      multi-homing;
      site-preference backup;
      interface ge-2/0/2.0 {
        remote-site-id 2;
      }
    }
  }
}
```

At PE3:

```
[edit routing-instances]
foo {
  instance-type l2vpn;
  interface ge-2/1/2.0;
  route-distinguisher 10.255.183.61:1;
  vrf-target target:9000:1;
  protocols {
    l2vpn {
      encapsulation-type ethernet-vlan;
      site foo {
        site-identifier 2;
        interface ge-2/1/2.0;
      }
    }
  }
}
```

Release History Table

Release	Description
14.2	Starting in Junos OS Release 14.2, Junos OS supports the restoration of egress traffic when there is a link or node failure in the egress PE node.

RELATED DOCUMENTATION

[Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services | 176](#)
site-preference

Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services

IN THIS SECTION

- [Requirements | 176](#)
- [Overview | 176](#)
- [Configuration | 178](#)
- [Verification | 194](#)

Starting in Junos OS Release 14.2, Junos OS supports the restoration of egress traffic when there is a link or node failure in the egress PE node. If there is a link or node failure in the core network, a protection mechanism such as MPLS fast reroute can be triggered on the transport LSPs between the PE routers to repair the connection within tens of milliseconds. An egress protection LSP addresses the problem of a node-link failure at the edge of the network (for example, a failure of a PE router).

This example shows how to configure link protection for BGP signaled Layer 2 services.

Requirements

MX Series Routers running Junos OS Release 14.2 or later.

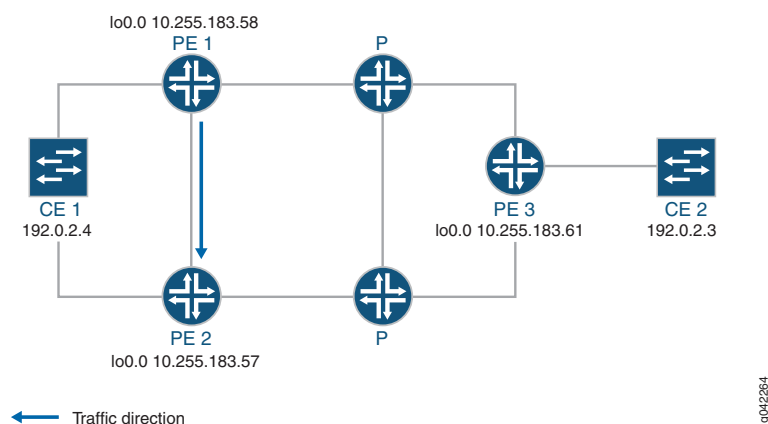
Overview

If there is a link or node failure in the core network, a protection mechanism such as MPLS fast reroute can be triggered on the transport LSPs between the PE routers to repair the connection within tens of milliseconds. An egress protection LSP addresses the problem of a node-link failure at the edge of the network (for example, a failure of a PE router).

This example includes the following configuration concepts and statements that are unique to the configuration of an egress protection LSP:

- **context-identifier**—Specifies an IPv4 or IPv6 address used to define the pair of PE routers participating in the egress protection LSP. It is assigned to each ordered pair of primary PE and the protector to facilitate protection establishment. This address is globally unique, or unique in the address space of the network where the primary PE and the protector reside.
- **egress-protection**—Configures the protector information for the protected Layer 2 circuit and configures the protector Layer 2 circuit at the `[edit protocols mpls]` hierarchy level. Configures an LSP as an egress protection LSP at the `[edit protocols mpls]` hierarchy level.
- **protector**—Configures the creation of standby pseudowires on the backup PE for link or node protection for the instance.

Figure 12: Egress Protection LSP Configured from Router PE1 to Router PE2



In the event of a failure of the egress PE Router PE1, traffic is switched to the egress protection LSP configured between Router PE1 and Router PE2 (the protector PE router):

- Device CE2—Traffic origin
- Router PE3—Ingress PE router
- Router PE1— (Primary) Egress PE router
- Router PE2—Protector PE router
- Device CE1—Traffic destination

When the link between CE1– PE1 goes down, PE1 will briefly redirect that traffic toward CE1, to PE2. PE2 forwards it to CE1 until ingress router PE3 recalculates to forward the traffic to PE2.

Initially the traffic direction was: CE2 – PE3 – P – PE1 – CE1.

When the link between CE1– PE1 goes down, the traffic will be: CE2 – PE3 – P – PE1 – PE2 –CE1. PE3 then recalculates the path: CE2 – PE3 – P – PE2 – CE1.

This example shows how to configure routers PE1, PE2, and PE3.

Configuration

IN THIS SECTION

- [Step-by-Step Procedure | 181](#)
- [Results | 187](#)

CLI Quick Configuration

To quickly configure an egress protection LSP, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configurations, copy and then paste the commands into the CLI and enter **commit** from configuration mode.

PE1

```
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 198.51.100.3 primary
set protocols mpls egress-protection context-identifier 198.51.100.3 advertise-mode stub-alias
set protocols mpls egress-protection traceoptions file ep size 100m
set protocols mpls egress-protection traceoptions flag all
set protocols bgp traceoptions file bgp.log world-readable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.183.58
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family l2vpn signaling egress-protection
set protocols bgp group ibgp neighbor 192.0.2.3
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols isis traceoptions file isis-edge size 10m world-readable
set protocols isis traceoptions flag error
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols ldp interface all
```



```

set protocols ldp interface fxp0.0 disable
set policy-options policy-statement lb then load-balance per-packet
set routing-options traceoptions file ro.log
set routing-options traceoptions flag all
set routing-options traceoptions flag route
set routing-options autonomous-system 100
set routing-options forwarding-table export lb
set routing-instances foo instance-type l2vpn
set routing-instances foo egress-protection context-identifier 198.51.100.3
set routing-instances foo interface ge-2/0/2.0
set routing-instances foo route-distinguisher 10.255.183.58:1
set routing-instances foo vrf-target target:9000:1
set routing-instances foo protocols l2vpn encapsulation-type ethernet-vlan
set routing-instances foo protocols l2vpn site foo site-identifier 1
set routing-instances foo protocols l2vpn site foo site-preference primary
set routing-instances foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2

```

PE2

```

set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 198.51.100.3 protector
set protocols mpls egress-protection context-identifier 198.51.100.3 advertise-mode stub-alias
set protocols mpls egress-protection traceoptions file ep size 100m
set protocols mpls egress-protection traceoptions flag all
set protocols bgp traceoptions file bgp.log world-readable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.183.57
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family l2vpn signaling egress-protection
set protocols bgp group ibgp neighbor 192.0.2.3
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols isis traceoptions file isis-edge size 10m world-readable
set protocols isis traceoptions flag error
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement lb then load-balance per-packet

```



```

set routing-options traceoptions file ro.log
set routing-options traceoptions flag normal
set routing-options traceoptions flag route
set routing-options autonomous-system 100
set routing-options forwarding-table export lb
set routing-instances foo instance-type l2vpn
set routing-instances foo egress-protection protector
set routing-instances foo interface ge-2/0/2.0
set routing-instances foo route-distinguisher 10.255.183.57:1
set routing-instances foo vrf-target target:9000:1
set routing-instances foo protocols l2vpn encapsulation-type ethernet-vlan
set routing-instances foo protocols l2vpn site foo hot-standby
set routing-instances foo protocols l2vpn site foo site-identifier 1
set routing-instances foo protocols l2vpn site foo site-preference backup
set routing-instances foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2

```

PE3

```

set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp traceoptions file bgp.log world-readable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.183.61
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family l2vpn signaling
set protocols bgp group ibgp neighbor 192.0.2.3
set protocols bgp group ibgp neighbor 192.0.2.4
set protocols isis traceoptions file isis-edge size 10m world-readable
set protocols isis traceoptions flag error
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement lb then load-balance per-packet
set routing-options traceoptions file ro.log
set routing-options traceoptions flag normal
set routing-options traceoptions flag route
set routing-options autonomous-system 100
set routing-options forwarding-table export lb

```



```

set routing-instances foo instance-type l2vpn
set routing-instances foo interface ge-2/1/2.0
set routing-instances foo route-distinguisher 10.255.183.61:1
set routing-instances foo vrf-target target:9000:1
set routing-instances foo protocols l2vpn encapsulation-type ethernet-vlan
set routing-instances foo protocols l2vpn site foo site-identifier 2
set routing-instances foo protocols l2vpn site foo interface ge-2/1/2.0 remote-site-id 1

```

Step-by-Step Procedure

Step-by-Step Procedure

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

To configure an egress protection LSP for router PE1:

1. Configure RSVP.

```

[edit protocols rsvp]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable

```

2. Configure MPLS to use the egress protection LSP to protect against a link failure to Device CE1.

```

[edit protocols mpls]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
user@PE1# set egress-protection context-identifier 198.51.100.3 primary
user@PE1# set egress-protection context-identifier 198.51.100.3 advertise-mode stub-alias
user@PE1# set egress-protection traceoptions file ep size 100m
user@PE1# set egress-protection traceoptions flag all

```

3. Configure BGP.

```

[edit protocols bgp]
user@PE1# set traceoptions file bgp.log world-readable
user@PE1# set group ibgp type internal
user@PE1# set group ibgp local-address 10.255.183.58
user@PE1# set group ibgp family inet unicast
user@PE1# set group ibgp family l2vpn signaling egress-protection
user@PE1# set group ibgp neighbor 192.0.2.3
user@PE1# set group ibgp neighbor 192.0.2.4

```


4. Configure IS-IS.

```
[edit protocols isis]
user@PE1# set traceoptions file isis-edge size 10m world-readable
user@PE1# set traceoptions flag error
user@PE1# set level 1 disable
user@PE1# set level 2 wide-metrics-only
user@PE1# set interface all point-to-point
user@PE1# set interface all level 2 metric 10
user@PE1# set interface fxp0.0 disable
```

5. Configure LDP.

```
[edit protocols ldp]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
```

6. Configure a load-balancing policy.

```
[edit]
user@PE1# set policy-options policy-statement lb then load-balance per-packet
```

7. Configure the routing options to export routes based on the load-balancing policy.

```
[edit routing-options]
user@PE1# set traceoptions file ro.log
user@PE1# set traceoptions flag all
user@PE1# set autonomous-system 100
user@PE1# set forwarding-table export lb
```

8. Configure BGP to advertise nrli from the routing instance with context-ID as next-hop.

```
[edit routing-instances]
user@PE1# set foo instance-type l2vpn
user@PE1# set foo egress-protection context-identifier 198.51.100.3
user@PE1# set foo interface ge-2/0/2.0
user@PE1# set foo route-distinguisher 10.255.183.58:1
user@PE1# set foo vrf-target target:9000:1
```

9. Configure l2vpn instance to use the egress LSP configured.


```
[edit routing-instances]
user@PE1# set foo protocols l2vpn encapsulation-type ethernet-vlan
user@PE1# set foo protocols l2vpn site foo site-identifier 1
user@PE1# set foo protocols l2vpn site foo site-preference primary
user@PE1# set foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2
```

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

Step-by-Step Procedure

To configure an egress protection LSP for Router PE2:

1. Configure RSVP.

```
[edit protocols rsvp]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
```

2. Configure MPLS and the LSP that acts as the egress protection LSP.

```
[edit protocols mpls]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
user@PE2# set egress-protection context-identifier 198.51.100.3 protector
user@PE2# set egress-protection context-identifier 198.51.100.3 advertise-mode stub-alias
user@PE2# set egress-protection traceoptions file ep size 100m
user@PE2# set egress-protection traceoptions flag all
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE2# set traceoptions file bgp.log world-readable
user@PE2# set group ibgp type internal
user@PE2# set group ibgp local-address 10.255.183.57
user@PE2# set group ibgp family inet unicast
user@PE2# set group ibgp family l2vpn signaling
user@PE2# set group ibgp family l2vpn egress-protection
user@PE2# set group ibgp neighbor 192.0.2.3
user@PE2# set group ibgp neighbor 192.0.2.4
```

4. Configure IS-IS.


```
[edit protocols isis]
user@PE2# set traceoptions file isis-edge size 10m world-readable
user@PE2# set traceoptions flag error
user@PE2# set level 1 disable
user@PE2# set level 2 wide-metrics-only
user@PE2# set interface all point-to-point
user@PE2# set interface all level 2 metric 10
user@PE2# set interface fxp0.0 disable
```

5. Configure LDP.

```
[edit protocols ldp]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
```

6. Configure a load-balancing policy.

```
[edit]
user@PE2# set policy-options policy-statement lb then load-balance per-packet
```

7. Configure the routing options to export routes based on the load-balancing policy.

```
[edit routing-options]
user@PE2# set traceoptions file ro.log
user@PE2# set traceoptions flag all
user@PE2# set autonomous-system 100
user@PE2# set forwarding-table export lb
```

8. Configure BGP to advertise nrli from the routing instance with context-ID as next-hop.

```
[edit routing-instances]
user@PE2# set foo instance-type l2vpn
user@PE2# set foo egress-protection protector
user@PE2# set foo interface ge-2/0/2.0
user@PE2# set foo route-distinguisher 10.255.183.57:1
user@PE2# set foo vrf-target target:9000:1
```

9. Configure l2vpn instance to use the egress LSP configured.


```
[edit routing-instances]
user@PE2# set foo protocols l2vpn encapsulation-type ethernet-vlan
user@PE2# set foo protocols l2vpn site foo hot-standby
user@PE2# set foo protocols l2vpn site foo site-identifier 1
user@PE2# set foo protocols l2vpn site foo site-preference backup
user@PE2# set foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2
```

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

Step-by-Step Procedure

To configure an egress protection LSP for Router PE3:

1. Configure RSVP.

```
[edit protocols rsvp]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
```

2. Configure MPLS.

```
[edit protocols mpls]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE3# set traceoptions file bgp.log world-readable
user@PE3# set group ibgp type internal
user@PE3# set group ibgp local-address 10.255.183.61
user@PE3# set group ibgp family inet unicast
user@PE3# set group ibgp family l2vpn signaling
user@PE3# set group ibgp neighbor 192.0.2.3
user@PE3# set group ibgp neighbor 192.0.2.4
```

4. Configure IS-IS.

```
[edit protocols isis]
user@PE3# set traceoptions file isis-edge size 10m world-readable
user@PE3# set traceoptions flag error
```



```

user@PE3# set level 1 disable
user@PE3# set level 2 wide-metrics-only
user@PE3# set protocols isis interface all point-to-point
[edit protocols isis]
user@PE3# set protocols isis interface all level 2 metric 10
[edit protocols isis]
user@PE3# set protocols isis interface fxp0.0 disable

```

5. Configure LDP.

```

[edit protocols ldp]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable

```

6. Configure a load-balancing policy.

```

[edit]
user@PE3# set policy-options policy-statement lb then load-balance per-packet

```

7. Configure the routing options to export routes based on the load-balancing policy.

```

[edit routing-options]
user@PE3# set traceoptions file ro.log
user@PE3# set traceoptions flag normal
user@PE3# set traceoptions flag route
user@PE3# set autonomous-system 100
user@PE3# set forwarding-table export lb

```

8. Configure BGP to advertise nlri from the routing instance with context-ID as next-hop.

```

[edit]
user@PE3# set routing-instances foo instance-type l2vpn
user@PE3# set routing-instances foo interface ge-2/1/2.0
user@PE3# set routing-instances foo route-distinguisher 10.255.183.61:1
user@PE3# set routing-instances foo vrf-target target:9000:1

```

9. Configure l2vpn to specify the interface that connects to the site and the remote interface to which you want the specified interface to connect.


```
[edit routing-instances]
user@PE3# set foo protocols l2vpn encapsulation-type ethernet-vlan
user@PE3# set foo protocols l2vpn site foo site-identifier 2
user@PE3# set foo protocols l2vpn site foo interface ge-2/1/2.0 remote-site-id 1
```

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

Results

From configuration mode, confirm your configuration on Router PE1 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@PE1# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 198.51.100.3 {
      primary;
      advertise-mode stub-alias;
    }
    traceoptions {
      file ep size 100m;
      flag all;
    }
  }
}
bgp {
  traceoptions {
    file bgp.log world-readable;
  }
  group ibgp {
```



```

    type internal;
    local-address 10.255.183.58;
    family inet {
        unicast;
    }
    family l2vpn {
        signaling {
            egress-protection;
        }
    }
    neighbor 192.0.2.3;
    neighbor 192.0.2.4;
}
}
isis {
    traceoptions {
        file isis-edge size 10m world-readable;
        flag error;
    }
    level 1 disable;
    level 2 wide-metrics-only;
    interface all {
        point-to-point;
        level 2 metric 10;
    }
    interface fxp0.0 {
        disable;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

[edit]
user@PE1# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}
[edit]

```



```

user@PE1# show routing-options
traceoptions {
  file ro.log;
  flag all;
}
autonomous-system 100;
forwarding-table {
  export lb;
}

[edit]
user@PE1# show routing-instances
foo {
  instance-type l2vpn;
  egress-protection {
    context-identifier {
      198.51.100.3;
    }
  }
}
interface ge-2/0/2.0;
route-distinguisher 10.255.183.58:1;
vrf-target target:9000:1;
protocols {
  l2vpn {
    encapsulation-type ethernet-vlan;
    site foo {
      site-identifier 1;
      site-preference primary;
      interface ge-2/0/2.0 {
        remote-site-id 2;
      }
    }
  }
}
}

```

From configuration mode, confirm your configuration on Router PE2 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit]
user@PE2# show protocols
rsvp {
  interface all;
}

```



```

    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
    egress-protection {
        context-identifier 198.51.100.3 {
            protector;
            advertise-mode stub-alias;
        }
        traceoptions {
            file ep size 100m;
            flag all;
        }
    }
}
bgp {
    traceoptions {
        file bgp.log world-readable;
    }
    group ibgp {
        type internal;
        local-address 10.255.183.57;
        family inet {
            unicast;
        }
        family l2vpn {
            signaling {
                egress-protection;
            }
        }
        neighbor 192.0.2.3;
        neighbor 192.0.2.4;
    }
}
isis {
    traceoptions {
        file isis-edge size 10m world-readable;
        flag error;
    }
}

```



```

level 1 disable;
level 2 wide-metrics-only;
interface all {
    point-to-point;
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```

```

[edit]
user@PE2# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}

```

```

[edit]
user@PE2# show routing-options
traceoptions {
    file ro.log;
    flag normal;
    flag route;
}
autonomous-system 100;
forwarding-table {
    export lb;
}

```

```

[edit]
user@PE2# show routing-instances
foo {
    instance-type l2vpn;
    egress-protection {
        protector;
    }
}

```



```

interface ge-2/0/2.0;
route-distinguisher 10.255.183.57:1;
vrf-target target:9000:1;
protocols {
  l2vpn {
    encapsulation-type ethernet-vlan;
    site foo {
      hot-standby;
      site-identifier 1;
      site-preference backup;
      interface ge-2/0/2.0 {
        remote-site-id 2;
      }
    }
  }
}

```

From configuration mode, confirm your configuration on Router PE3 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit]
user@PE3# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  traceoptions {
    file bgp.log world-readable;
  }
  group ibgp {
    type internal;
    local-address 10.255.183.61;
    family inet {

```



```

        unicast;
    }
    family l2vpn {
        signaling;
    }
    neighbor 192.0.2.3;
    neighbor 192.0.2.4;
    }
}
isis {
    traceoptions {
        file isis-edge size 10m world-readable;
        flag error;
    }
    level 1 disable;
    level 2 wide-metrics-only;
    interface all {
        point-to-point;
        level 2 metric 10;
    }
    interface fxp0.0 {
        disable;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```

[edit]

user@PE3# **show policy-options**

```

policy-statement lb {
    then {
        load-balance per-packet;
    }
}

```

[edit]

user@PE3# **show routing-options**

```

traceoptions {
    file ro.log;
    flag normal;
}

```



```

    flag route;
}
autonomous-system 100;
forwarding-table {
    export lb;
}

[edit]
user@PE3# show routing-instances
foo {
    instance-type l2vpn;
    interface ge-2/1/2.0;
    route-distinguisher 10.255.183.61:1;
    vrf-target target:9000:1;
    protocols {
        l2vpn {
            encapsulation-type ethernet-vlan;
            site foo {
                site-identifier 2;
                interface ge-2/1/2.0 {
                    remote-site-id 1;
                }
            }
        }
    }
}

```

Verification

IN THIS SECTION

- [Verifying the L2VPN Configuration | 195](#)
- [Verifying the Routing Instance Details | 196](#)
- [Verifying the IS-IS Configuration | 196](#)
- [Verifying the MPLS Configuration | 197](#)

Confirm that the configuration is working properly.

Verifying the L2VPN Configuration

Purpose

Verify that LSP is protected by the connection protection logic.

Action

From operational mode, run the **show l2vpn connections extensive** command.

```
user@PE2> show l2vpn connections extensive
```

```
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label         Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch              MI -- Mesh-Group ID not available
BK -- Backup connection         ST -- Standby connection
PF -- Profile parse failure     PB -- Profile busy
RS -- remote site standby       SN -- Static Neighbor
LB -- Local site not best-site  RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down

Instance: foo
Local site: foo (1)
  connection-site      Type  St  Time last up      # Up trans
    2                  rmt   Up   Aug  3 00:08:14 2001      1
    Local circuit: ge-2/0/2.0, Status: Up
    Remote PE: 192.0.2.3
    Incoming label: 32769, Outgoing label: 32768
    Egress Protection: Yes
      Time              Event              Interface/Lbl/PE
    Aug  3 00:08:14 2001 PE route up
```



```

Aug  3 00:08:14 2001  Out lbl Update          32768
Aug  3 00:08:14 2001  In  lbl Update          32769
Aug  3 00:08:14 2001  ckt0 up              fe-0/0/0.0

```

Meaning

The **Egress Protection: Yes** output shows that the given PVC is protected by connection protection logic.

Verifying the Routing Instance Details

Purpose

Verify the routing instance information and the context identifier configured on the primary, which is used as the next-hop address in case of node-link failure.

Action

From operational mode, run the **show route foo detail** command.

```
user@PE2> show route foo detail
```

```

foo:
  Router ID: 0.0.0.0
  Type: l2vpn non-forwarding State: Active
  Interfaces:
    lt-1/2/0.56
  Route-distinguisher: 10.255.255.11:1
  Vrf-import: [ __vrf-import-foo-internal__ ]
  Vrf-export: [ __vrf-export-foo-internal__ ]
  Vrf-import-target: [ target:100:200 ]
  Vrf-export-target: [ target:100:200 ]
  Fast-reroute-priority: low
  Vrf-edge-protection-id: 198.51.100.3
  Tables:
    foo.l2vpn.0          : 5 routes (3 active, 0 holddown, 0 hidden)
    foo.l2id.0           : 6 routes (2 active, 0 holddown, 0 hidden)

```

Meaning

The context-id is set to **198.51.100.3** and the **Vrf-import: [__vrf-import-foo-internal__]** in the output mentions the policy used for rewriting the next-hop address.

Verifying the IS-IS Configuration

Purpose

Verify the IS-IS context identifier information.

Action

From operational mode, run the **show isis context-identifier detail** command.

```
user@PE2> show isis context-identifier detail
```

```
IS-IS context database:
```

Context	L	Owner	Role	Primary	Metric
198.51.100.3	2	MPLS	Protector	pro17-b-lr-R1	0
Advertiser pro17-b, Router ID 10.255.107.49, Level 2, tlv protector					
Advertiser pro17-b-lr-R1, Router ID 10.255.255.11, Metric 1, Level 2, tlv prefix					

Meaning

Router PE2 is the protector and the configured context identifier is in use for the MPLS protocol.

Verifying the MPLS Configuration

Purpose

Verify the context identifier details on the primary and protector PEs.

Action

From operational mode, run the **show mpls context-identifier detail** command.

```
user@PE1> show mpls context-identifier detail
```

```
ID: 198.51.100.3
  Type: primary, Metric: 1, Mode: alias

Total 1, Primary 1, Protector 0
```

```
user@PE2> show mpls context-identifier detail
```

```
ID: 198.51.100.3
  Type: protector, Metric: 16777215, Mode: alias
  Context table: __198.51.100.3__.mpls.0, Label out: 299968
```



```
user@PE2> show mpls egress-protection detail
```

```
Instance          Type      Protection-Type
foo               local-l2vpn Protector
Route Target 100:200
```

Meaning

Context-id is **198.51.100.3**, advertise-mode is **alias**, the MPLS table created for egress protection is **__198.51.100.3__.mpls.0**, and the egress instance name is **foo**, which is of type **local-l2vpn**.

Release History Table

Release	Description
14.2	Starting in Junos OS Release 14.2, Junos OS supports the restoration of egress traffic when there is a link or node failure in the egress PE node.

RELATED DOCUMENTATION

Configuring Per-Packet Load Balancing

Introduction to Configuring Layer 2 VPNs

site-preference

Example: Configuring Layer 3 VPN Egress Protection with PLR as Protector

IN THIS SECTION

- [Requirements | 199](#)
- [Overview | 199](#)
- [Configuration | 200](#)
- [Verification | 223](#)

This example shows how to configure fast service restoration at the egress of a Layer 3 VPN when the customer is multihomed to the service provider.

Starting in Junos OS Release 15.1, the enhanced point of local repair (PLR) functionality addresses a special scenario of egress node protection, where the PLR and the protector are co-located as one router. In this case, there is no need to have a bypass LSP reroute traffic during local repair. Instead, the PLR or the protector can send the traffic directly to the target CE (in Co-located protector model where the PLR or the protector is also the backup PE that is directly connected to the CE) or to the backup PE (in Centralized protector model where the backup PE is a separate router).

Requirements

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

This example requires Junos OS Release 15.1 or later.

Overview


As a special scenario of egress node protection, if a router is both a Protector and a PLR, it installs backup next hops to protect the transport LSP. In particular, it does not need a bypass LSP for local repair.

In the Co-located protector model, the PLR or the Protector is directly connected to the CE via a backup AC, while in the Centralized protector model, the PLR or the protector has an MPLS tunnel to the backup PE. In either case, the PLR or the Protector will install a backup next hop with a label followed by a lookup in a **context label** table, i.e. `__context__.mpls.0`. When the egress node fails, the PLR or the Protector will switch traffic to this backup next hop in PFE. The outer label (the transport LSP label) of packets is popped, and the inner label (the layer 3 VPN label allocated by the egress node) is looked up in `__context__.mpls.0`, which results in forwarding the packets directly to the CE (in Collocated protector model) or the backup PE (in Centralized protector model).

Topology

[Figure 13 on page 200](#) shows the sample network.

Figure 13: Co-located PLR and protector in collocated protector model



Configuration

IN THIS SECTION

- [Configuring Device CE1 | 205](#)
- [Configuring Device PE1 | 205](#)
- [Configuring Device P | 208](#)
- [Configuring Device PE2 | 209](#)
- [Configuring Device PE3 | 211](#)
- [Configuring Device CE2 | 213](#)
- [Results | 213](#)

CLI Quick Configuration

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

Device CE1

```
set interfaces ge-0/0/0 unit 0 family inet address 10.10.20.2/30
set interfaces lo0 unit 0 family inet address 10.255.162.87/32
```

Device PE1


```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.20.1/30
set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.1/30
set interfaces ge-0/0/1 unit 0 family inet6
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.84/32 primary
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:84/128 primary
set interfaces lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2084.00
set policy-options policy-statement vpn-exp term 1 from protocol direct
set policy-options policy-statement vpn-exp term 1 from route filter 10.10.20.0/24 exact
set policy-options policy-statement vpn-exp term 1 then community add vpn
set policy-options policy-statement vpn-exp term 1 then accept
set policy-options policy-statement vpn-imp term 1 from community vpn
set policy-options policy-statement vpn-imp term 1 then accept
set policy-options policy-statement vpn-imp term 2 then reject
set policy-options community vpn members traget:1:1
set routing-options autonomous-system 65000
set protocols rsvp interface all link-protection
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp vpn-apply-export
set protocols bgp group vpn type internal
set protocols bgp group vpn local-address 10.255.162.84
set protocols bgp group vpn family inet-vpn unicast
set protocols bgp group vpn neighbor 10.255.162.91
set protocols bgp group vpn neighbor 10.255.162.89
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive
set routing-instances vpn instance-type vrf
set routing-instances vpn interface ge-1/0/0.0
set routing-instances vpn route-distinguisher 100:100
set routing-instances vpn vrf-import vpn-imp
set routing-instances vpn vrf-export vpn-exp
set routing-instances vpn vrf-table-label
set routing-instances vpn protocols bgp group vpn type external
set routing-instances vpn protocols bgp group vpn family inet unicast
set routing-instances vpn protocols bgp group vpn family inet6 unicast
set routing-instances vpn protocols bgp group vpn peer-as 65001
set routing-instances vpn protocols bgp group vpn as-override

```



```
set routing-instances vpn protocols bgp group vpn neighbor 10.10.20.2
```

Device P

```
set interfaces ge-0/0/0 unit 0 family inet address 10.10.11.2/30
set interfaces ge-0/0/0 unit 0 family inet6
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.2/30
set interfaces ge-0/0/1 unit 0 family inet6
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.86/32 primary
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:86/128 primary
set interfaces lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2086.00
set protocols rsvp interface all link-protection
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
```

Device PE2

```
set interfaces ge-0/0/0 unit 0 family inet address 10.10.11.1/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.12.1/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.10.30.1/30
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.91/32 primary
set interfaces lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2091.00
```



```
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:91/128 primary
set routing-options graceful-restart
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols rsvp interface all link-protection
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path to_PE1 to 10.255.162.84
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 1.1.1.1 protector
set protocols mpls egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias
set protocols bgp vpn-apply-export
set protocols bgp group vpn type internal
set protocols bgp group vpn local-address 10.255.162.91
set protocols bgp group vpn family inet-vpn unicast egress-protection
set protocols bgp group vpn neighbor 10.255.162.84
set protocols bgp group vpn neighbor 10.255.162.89
set protocols isis traceoptions file isis.log
set protocols isis traceoptions flag all detail
set protocols isis level 2 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive
set policy-options policy-statement pplb term 1 then load-balance per-packet
set policy-options policy-statement vpn-exp term 1 from protocol bgp
set policy-options policy-statement vpn-exp term 1 then community add vpn
set policy-options policy-statement vpn-exp term 1 then accept
set policy-options policy-statement vpn-imp term 1 from community vpn
set policy-options policy-statement vpn-imp term 1 then accept
set policy-options policy-statement vpn-imp term 2 then reject
set policy-options community vpn members target:1:1
set routing-instances vpn instance-type vrf
set routing-instances vpn interface ge-3/2/4.0
set routing-instances vpn route-distinguisher 100:100
set routing-instances vpn vrf-import vpn-imp
set routing-instances vpn vrf-export vpn-exp
set routing-instances vpn vrf-table-label
set routing-instances vpn protocols bgp group vpn type external
set routing-instances vpn protocols bgp group vpn family inet unicast
set routing-instances vpn protocols bgp group vpn family inet6 unicast
set routing-instances vpn protocols bgp group vpn peer-as 65001
set routing-instances vpn protocols bgp group vpn as-override
```



```
set routing-instances vpn protocols bgp group vpn neighbor 10.10.30.2
```

Device PE3

```
set interfaces ge-0/0/0 unit 0 family inet address 10.10.40.1/30
set interfaces ge-0/0/1 unit 0 family inet address 10.10.12.2/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.89/32 primary
set interfaces lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2089.00
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:89/128 primary
set routing-options graceful-restart
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols rsvp interface all link-protection
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path to_PE2 to 10.255.162.91
set protocols mpls label-switched-path to_PE1 to 10.255.162.84
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 1.1.1.1 primary
set protocols mpls egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias
set protocols bgp vpn-apply-export
set protocols bgp group vpn type internal
set protocols bgp group vpn local-address 10.255.162.89
set protocols bgp group vpn family inet-vpn unicast
set protocols bgp group vpn neighbor 10.255.162.84 local-preference 300
set protocols bgp group vpn neighbor 10.255.162.91
set protocols isis level 2 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive
set routing-instances vpn instance-type vrf
set routing-instances vpn egress-protection context-identifier 1.1.1.1
set routing-instances vpn interface ge-1/1/0.0
set routing-instances vpn route-distinguisher 100:100
set routing-instances vpn vrf-import vpn-imp
set routing-instances vpn vrf-export vpn-exp
```



```

set routing-instances vpn vrf-table-label
set routing-instances vpn protocols bgp group vpn type external
set routing-instances vpn protocols bgp group vpn family inet unicast
set routing-instances vpn protocols bgp group vpn family inet6 unicast
set routing-instances vpn protocols bgp group vpn peer-as 65001
set routing-instances vpn protocols bgp group vpn as-override
set routing-instances vpn protocols bgp group vpn neighbor 10.10.40.2

```

Device CE2

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.40.2/30
set interfaces ge-0/0/2 unit 0 family inet address 10.10.30.2/30
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.88/32 primary
set interfaces lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2088.00
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:88/128 primary

```

Configuring Device CE1

Step-by-Step Procedure

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

1. Configure interfaces.

```

[edit interfaces]
user@CE1# set ge-0/0/0 unit 0 family inet address 10.10.20.2/30
user@CE1# set lo0 unit 0 family inet address 10.255.162.87/32

```

Configuring Device PE1

Step-by-Step Procedure

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 10.10.20.1/30
user@PE1# set ge-0/0/1 unit 0 family inet address 10.10.10.1/30
user@PE1# set ge-0/0/1 unit 0 family iso

```



```

user@PE1# set ge-0/0/1 unit 0 family inet6
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 127.0.0.1/32
user@PE1# set lo0 unit 0 family inet address 10.255.162.84/32 primary
user@PE1# set lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2084.00
user@PE1# set lo0 unit 0 family inet6 address abcd::10:255:162:84/128 primary

```

2. Configure the autonomous system (AS) number.

```

[edit routing-options]
user@PE1# set autonomous-system 65000
user@PE1# set forwarding-table export pplb

```

3. Configure RSVP.

```

[edit protocols rsvp]
user@PE1# set interface all link-protection
user@PE1# set interface fxp0.0 disable

```

4. Enable MPLS.

```

[edit protocols mpls]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable

```

5. Configure BGP.

```

[edit protocols bgp]
user@PE1# set group vpn type internal
user@PE1# set group vpn local-address 10.255.162.84
user@PE1# set group vpn family inet-vpn unicast
user@PE1# set group vpn neighbor 10.255.162.91
user@PE1# set group vpn neighbor 10.255.162.89
user@PE1# set vpn-apply-export

```

6. Enable IS-IS.

```

[edit protocols isis]
user@PE1# set interface all

```



```

user@PE1# set interface fxp0.0 disable
user@PE1# set interface lo0.0 passive

```

7. (Optional) Configure OSPF

```

[edit protocols ospf]
user@PE1# set area 0.0.0.0 interface all
user@PE1# set area 0.0.0.0 interface fxp0.0 disable
user@PE1# set area 0.0.0.0 interface lo0.0 passive
user@PE1# set traffic-engineering

```

8. Configure the routing instance.

```

[edit routing-instances]
user@PE1# set vpn instance-type vrf
user@PE1# set vpn interface ge-1/0/0.0
user@PE1# set vpn route-distinguisher 100:100
user@PE1# set vpn vrf-import vpn-imp
user@PE1# set vpn vrf-export vpn-exp
user@PE1# set vpn vrf-table-label
user@PE1# set vpn protocols bgp group vpn type external
user@PE1# set vpn protocols bgp group vpn family inet unicast
user@PE1# set vpn protocols bgp group vpn family inet6 unicast
user@PE1# set vpn protocols bgp group vpn peer-as 65001
user@PE1# set vpn protocols bgp group vpn as-override
user@PE1# set vpn protocols bgp group vpn neighbor 10.10.20.2

```

9. Configure the routing policy.

```

[edit]
user@PE1# set policy-options policy-statement vpn-exp term 1 from protocol direct
user@PE1# set policy-options policy-statement vpn-exp term 1 from route filter 10.10.20.0/24 exact
user@PE1# set policy-options policy-statement vpn-exp term 1 then community add vpn
user@PE1# set policy-options policy-statement vpn-exp term 1 then accept
user@PE1# set policy-options policy-statement vpn-imp term 1 from community vpn
user@PE1# set policy-options policy-statement vpn-imp term 1 then accept
user@PE1# set policy-options policy-statement vpn-imp term 2 then reject
user@PE1# set policy-options community vpn members traget:1:1

```


Configuring Device P

Step-by-Step Procedure

1. Configure the device interfaces.

```
[edit interfaces]
user@P# set ge-0/0/0 unit 0 family inet address 10.10.11.2/30
user@P# set ge-0/0/0 unit 0 family inet6
user@P# set ge-0/0/0 unit 0 family iso
user@P# set ge-0/0/0 unit 0 family mpls
user@P# set ge-0/0/1 unit 0 family inet address 10.10.10.2/30
user@P# set ge-0/0/1 unit 0 family inet6
user@P# set ge-0/0/1 unit 0 family iso
user@P# set ge-0/0/1 unit 0 family mpls
user@P# set lo0 unit 0 family inet address 127.0.0.1/32
user@P# set lo0 unit 0 family inet address 10.255.162.86/32 primary
user@P# set lo0 unit 0 family inet6 address abcd::10:255:162:86/128 primary
user@P# set lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2086.00
```

2. Enable IS-IS.

```
[edit protocols isis]
user@P# set interface all
user@P# set interface fxp0.0 disable
```

3. Enable MPLS.

```
[edit protocols mpls ]
user@P# set interface all
user@P# set interface fxp0.0 disable
```

4. Configure RSVP.

```
[edit protocols rsvp]
user@P# set interface all link-protection
user@P# set interface fxp0.0 disable
```

5. (Optional) Configure OSPF.

```
[edit protocols ospf]
```



```

user@P# set area 0.0.0.0 interface all
user@P# set area 0.0.0.0 interface fxp0.0 disable
user@P# set area 0.0.0.0 interface lo0.0 passive
user@P# set traffic-engineering

```

Configuring Device PE2

Step-by-Step Procedure

1. Configure the interfaces.

```

[edit interfaces]
user@PE2# set ge-0/0/0 unit 0 family inet address 10.10.11.1/30
user@PE2# set ge-0/0/0 unit 0 family iso
user@PE2# set ge-0/0/0 unit 0 family inet6
user@PE2# set ge-0/0/0 unit 0 family mpls
user@PE2# set ge-0/0/1 unit 0 family inet address 10.10.12.1/30
user@PE2# set ge-0/0/1 unit 0 family iso
user@PE2# set ge-0/0/1 unit 0 family inet6
user@PE2# set ge-0/0/1 unit 0 family mpls
user@PE2# set ge-0/0/2 unit 0 family inet address 10.10.30.1/30
user@PE2# set lo0 unit 0 family inet address 127.0.0.1/32
user@PE2# set lo0 unit 0 family inet address 10.255.162.91/32 primary
user@PE2# set lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2091.00
user@PE2# set lo0 unit 0 family inet6 address abcd::10:255:162:91/128 primary

```

2. Configure autonomous number(AS).

```

[edit routing-options]
user@PE2# set autonomous-system 65000
user@PE2# set forwarding-table export pplb

```

3. Configure RSVP.

```

[edit protocols rsvp]
user@PE2# set interface all link-protection
user@PE2# set interface fxp0.0 disable

```

4. Configure MPLS.

```

[edit protocols mpls]

```



```

user@PE2# set label-switched-path to_PE1 to 10.255.162.84
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
user@PE2# set egress-protection context-identifier 1.1.1.1 protector
user@PE2# set egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias

```

5. Configure BGP.

```

[edit protocols bgp]
user@PE2# set group vpn family inet-vpn unicast egress-protection
user@PE2# set group vpn local-address 10.255.162.91
user@PE2# set group vpn neighbor 10.255.162.84
user@PE2# set group vpn neighbor 10.255.162.89
user@PE2# set group vpn type internal
user@PE2# set vpn-apply-export

```

6. Configure IS-IS.

```

[edit protocols isis]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
user@PE2# set interface lo0.0 passive
user@PE2# set level 2 disable
user@PE2# set traceoptions file isis.log
user@PE2# set traceoptions flag all detail

```

7. (Optional) Configure OSPF.

```

[edit protocols ospf]
user@PE2# set area 0.0.0.0 interface all
user@PE2# set area 0.0.0.0 interface fxp0.0 disable
user@PE2# set area 0.0.0.0 interface lo0.0 passive
user@PE2# set traffic-engineering

```

8. Configure the routing policy.

```

[edit policy-options]
user@PE2# set community vpn members target:1:1
user@PE2# set policy-statement pplb term 1 then load-balance per-packet
user@PE2# set policy-statement vpn-exp term 1 from protocol bgp

```



```

user@PE2# set policy-statement vpn-exp term 1 then community add vpn
user@PE2# set policy-statement vpn-exp term 1 then accept
user@PE2# set policy-statement vpn-imp term 1 from community vpn
user@PE2# set policy-statement vpn-imp term 1 then accept
user@PE2# set policy-statement vpn-imp term 2 then reject

```

9. Configure the routing instance.

```

[edit routing-instances]
user@PE2# set vpn instance-type vrf
user@PE2# set vpn interface ge-3/2/4.0
user@PE2# set vpn route-distinguisher 100:100
user@PE2# set vpn vrf-import vpn-imp
user@PE2# set vpn vrf-export vpn-exp
user@PE2# set vpn vrf-table-label
user@PE2# set vpn protocols bgp group vpn type external
user@PE2# set vpn protocols bgp group vpn family inet unicast
user@PE2# set vpn protocols bgp group vpn family inet6 unicast
user@PE2# set vpn protocols bgp group vpn peer-as 65001
user@PE2# set vpn protocols bgp group vpn as-override
user@PE2# set vpn protocols bgp group vpn neighbor 10.10.30.2

```

Configuring Device PE3

Step-by-Step Procedure

1. Configure the interfaces.

```

[edit interfaces]
user@PE3# set ge-0/0/0 unit 0 family inet address 10.10.40.1/30
user@PE3# set ge-0/0/1 unit 0 family inet address 10.10.12.2/30
user@PE3# set ge-0/0/1 unit 0 family iso
user@PE3# set ge-0/0/1 unit 0 family inet6
user@PE3# set ge-0/0/1 unit 0 family mpls
user@PE3# set lo0 unit 0 family inet address 127.0.0.1/32
user@PE3# set lo0 unit 0 family inet address 10.255.162.89/32 primary
user@PE3# set lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2089.00
user@PE3# set lo0 unit 0 family inet6 address abcd::10:255:162:89/128 primary

```

2. Configure the autonomous number (AS).

```

[edit routing-options]

```



```
user@PE3# set autonomous-system 65000
user@PE3# set forwarding-table export pplb
```

3. Configure RSVP.

```
[edit protocols rsvp]
user@PE3# set interface all link-protection
user@PE3# set interface fxp0.0 disable
```

4. Configure MPLS.

```
[edit protocols mpls]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
user@PE3# set egress-protection context-identifier 1.1.1.1 primary
user@PE3# set egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias
user@PE3# set label-switched-path to_PE2 to 10.255.162.91
user@PE3# set label-switched-path to_PE1 to 10.255.162.84
```

5. Configure BGP.

```
[edit protocols bgp]
user@PE3# set group vpn type internal
user@PE3# set group vpn local-address 10.255.162.89
user@PE3# set group vpn family inet-vpn unicast
user@PE3# set group vpn neighbor 10.255.162.84 local-preference 300
user@PE3# set group vpn neighbor 10.255.162.91
user@PE3# set vpn-apply-export
```

6. Configure IS-IS.

```
[edit protocols isis]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
user@PE3# set interface lo0.0 passive
user@PE3# set level 2 disable
```

7. (Optional) Configure OSPF.


```
[edit protocols ospf]
user@PE3# set area 0.0.0.0 interface all
user@PE3# set area 0.0.0.0 interface fxp0.0 disable
user@PE3# set area 0.0.0.0 interface lo0.0 passive
user@PE3# set traffic-engineering
```

8. Configure the routing instance.

```
[edit routing-instances]
user@PE3# set vpn egress-protection context-identifier 1.1.1.1
user@PE3# set vpn instance-type vrf
user@PE3# set vpn interface ge-1/1/0.0
user@PE3# set vpn protocols bgp group vpn type external
user@PE3# set vpn protocols bgp group vpn family inet unicast
user@PE3# set vpn protocols bgp group vpn family inet6 unicast
user@PE3# set vpn protocols bgp group vpn peer-as 65001
user@PE3# set vpn protocols bgp group vpn as-override
user@PE3# set vpn protocols bgp group vpn neighbor 10.10.40.2
user@PE3# set vpn route-distinguisher 100:100
user@PE3# set vpn vrf-export vpn-exp
user@PE3# set vpn vrf-import vpn-imp
user@PE3# set vpn vrf-table-label
```

Configuring Device CE2

Step-by-Step Procedure

1. Configure the interfaces.

```
[edit interfaces]
user@CE2# set ge-0/0/0 unit 0 family inet address 10.10.40.2/30
user@CE2# set ge-0/0/2 unit 0 family inet address 10.10.30.2/30
user@CE2# set lo0 unit 0 family inet address 127.0.0.1/32
user@CE2# set lo0 unit 0 family inet address 10.255.162.88/32 primary
user@CE2# set lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2088.00
user@CE2# set lo0 unit 0 family inet6 address abcd::10:255:162:88/128 primary
```

Results

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

Device CE1

```
user@CE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.2/30;
    }
  }
}
```

Device PE1

```
user@PE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.1/30;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.10.10.1/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.162.84/32 {
        primary;
      }
    }
  }
  family iso {
```



```

        address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2084.00;
    }
    family inet6 {
        address abcd::10:255:162:84/128 {
            primary;
        }
    }
}
}

```

user@PE1# **show protocols**

```

rsvp {
    interface all {
        link-protection;
    }
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    vpn-apply-export;
    group vpn {
        type internal;
        local-address 10.255.162.84;
        family inet-vpn {
            unicast;
        }
        neighbor 10.255.162.91;
        neighbor 10.255.162.89;
    }
}
isis {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```



```

    }
    interface lo0.0 {
        passive;
    }
}

```

Device P

```

user@P# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.10.11.2/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 10.10.10.2/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            address 10.255.162.86/32 {
                primary;
            }
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2086.00;
        }
        family inet6 {
            address abcd::10:255:162:86/128 {

```



```

        primary;
    }
}
}
}

```

user@P# **show protocols**

```

rsvp {
  interface all {
    link-protection;
  }
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
isis {
  interface all;
  interface fxp0.0 {
    disable;
  }
}

```

Device PE2

user@PE2# **show interfaces**

```

ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.11.1/30;
    }
    family iso;
    family inet6;
  }
}

```



```

        family mpls;
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 10.10.12.1/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 10.10.30.1/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            address 10.255.162.91/32 {
                primary;
            }
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2091.00;
        }
        family inet6 {
            address abcd::10:255:162:91/128 {
                primary;
            }
        }
    }
}

```

```
user@PE2# show protocols
```

```
  rsvp {
```



```

interface all {
    link-protection;
}
interface fxp0.0 {
    disable;
}
}
mpls {
    label-switched-path to_PE1 {
        to 10.255.162.84;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
    egress-protection {
        context-identifier 1.1.1.1 {
            protector;
            advertise-mode stub-alias;
        }
    }
}
}
bgp {
    vpn-apply-export;
    group vpn {
        type internal;
        local-address 10.255.162.91;
        family inet-vpn {
            unicast {
                egress-protection;
            }
        }
        neighbor 10.255.162.84;
        neighbor 10.255.162.89;
    }
}
}
isis {
    traceoptions {
        file isis.log;
        flag all detail;
    }
    level 2 disable;
    interface all;
    interface fxp0.0 {

```



```

        disable;
    }
    interface lo0.0 {
        passive;
    }
}

```

Device PE3

```

user@PE3# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.10.40.1/30;
        }
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 10.10.12.2/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            address 10.255.162.89/32 {
                primary;
            }
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2089.00;
        }
        family inet6 {
            address abcd::10:255:162:89/128 {
                primary;
            }
        }
    }
}

```



```

    }
  }
}

```

```

user@PE3# show protocols
rsvp {
  interface all {
    link-protection;
  }
  interface fxp0.0 {
    disable;
  }
}
mpls {
  label-switched-path to_PE2 {
    to 10.255.162.91;
  }
  label-switched-path to_PE1 {
    to 10.255.162.84;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 1.1.1.1 {
      primary;
      advertise-mode stub-alias;
    }
  }
}
bgp {
  vpn-apply-export;
  group vpn {
    type internal;
    local-address 10.255.162.89;
    family inet-vpn {
      unicast;
    }
  }
  neighbor 10.255.162.84 {
    local-preference 300;
  }
}

```



```

    }
    neighbor 10.255.162.91;
  }
}
isis {
  level 2 disable;
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0 {
    passive;
  }
}

```

Device CE2

```

user@CE2# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.40.2/30;
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 10.10.30.2/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.162.88/32 {
        primary;
      }
    }
  }
  family iso {
    address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2088.00;
  }
}

```



```

    }
    family inet6 {
        address abcd::10:255:162:88/128 {
            primary;
        }
    }
}
}

```

Verification

IN THIS SECTION

- [Verifying the Routing Instance | 223](#)
- [Checking the Context Identifier Route | 231](#)

Verifying the Routing Instance

Purpose

Check the routes in the routing table.

Action

```
user@PE1> show route 10.10.50 table vpn.inet.0
```

```

vpn.inet.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.50.0/24      *[BGP/170] 00:01:26, localpref 100, from 10.255.162.96
                   AS path: 65001 I, validation-state: unverified
                   > to 10.10.10.2 via ge-2/0/2.0, Push 16, Push 300064(top)
                   [BGP/170] 00:06:22, localpref 50, from 10.255.162.91
                   AS path: 65001 I, validation-state: unverified
                   > to 10.10.10.2 via ge-2/0/2.0, Push 17, Push 299920(top)

```

```
user@PE1>show route 10.10.50 extensive table vpn.inet.0
```



```

vpn.inet.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
10.10.50.0/24 (2 entries, 1 announced)
TSI:
KRT in-kernel 10.10.50.0/24 -> {indirect(1048575)}
Page 0 idx 1, (group vpn type External) Type 1 val 0x9e33490 (adv_entry)
  Advertised metrics:
    Nexthop: Self
    AS path: [65000] 65000 I
    Communities: target:1:1
Path 10.10.50.0 from 10.255.162.96 Vector len 4. Val: 1
  *BGP      Preference: 170/-101
            Route Distinguisher: 200:100
            Next hop type: Indirect, Next hop index: 0
            Address: 0x9db63f0
            Next-hop reference count: 6
            Source: 10.255.162.96
            Next hop type: Router, Next hop index: 635
            Next hop: 10.10.10.2 via ge-2/0/2.0, selected
            Label operation: Push 16, Push 300064(top)
            Label TTL action: prop-ttl, prop-ttl(top)
            Load balance label: Label 16: None; Label 300064: None;
            Label element ptr: 0x9db60e0
            Label parent element ptr: 0x9db5e40
            Label element references: 1
            Label element child references: 0
            Label element lsp id: 0
            Session Id: 0x146
            Protocol next hop: 1.1.1.1
            Label operation: Push 16
            Label TTL action: prop-ttl
            Load balance label: Label 16: None;
            Indirect next hop: 0x9e55440 1048575 INH Session ID: 0x14d
            State: < Secondary Active Int Ext ProtectionCand >
            Local AS: 65000 Peer AS: 65000
            Age: 1:28 Metric2: 1
            Validation State: unverified
            Task: BGP_65000.10.255.162.96
            Announcement bits (2): 0-KRT 1-BGP_RT_Background
            AS path: 65001 I
            Communities: target:1:1
            Import Accepted
            VPN Label: 16
            Localpref: 100

```



```

Router ID: 10.255.162.96
Primary Routing Table bgp.l3vpn.0
Indirect next hops: 1
    Protocol next hop: 1.1.1.1 Metric: 1
    Label operation: Push 16
    Label TTL action: prop-ttl
    Load balance label: Label 16: None;
    Indirect next hop: 0x9e55440 1048575 INH Session ID: 0x14d

    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.10.10.2 via ge-2/0/2.0
        Session Id: 0x146
1.1.1.1/32 Originating RIB: inet.3
    Metric: 1      Node path count: 1
    Forwarding nexthops: 1
    Nexthop: 10.10.10.2 via ge-2/0/2.0
BGP      Preference: 170/-51
    Route Distinguisher: 100:100
    Next hop type: Indirect, Next hop index: 0
    Address: 0x9db6390
    Next-hop reference count: 5
    Source: 10.255.162.91
    Next hop type: Router, Next hop index: 636
    Next hop: 10.10.10.2 via ge-2/0/2.0, selected
    Label operation: Push 17, Push 299920(top)
    Label TTL action: prop-ttl, prop-ttl(top)
    Load balance label: Label 17: None; Label 299920: None;
    Label element ptr: 0x9db62c0
    Label parent element ptr: 0x9dc0d00
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
    Session Id: 0x146
    Protocol next hop: 10.255.162.91
    Label operation: Push 17
    Label TTL action: prop-ttl
    Load balance label: Label 17: None;
    Indirect next hop: 0x9e55580 1048574 INH Session ID: 0x14c
    State: < Secondary Int Ext ProtectionCand >
    Inactive reason: Local Preference
    Local AS: 65000 Peer AS: 65000
    Age: 6:24 Metric2: 1
    Validation State: unverified

```



```

Task: BGP_65000.10.255.162.91
AS path: 65001 I
Communities: target:1:1
Import Accepted
VPN Label: 17
Localpref: 50
Router ID: 10.255.162.91
Primary Routing Table bgp.l3vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.162.91 Metric: 1
    Label operation: Push 17
    Label TTL action: prop-ttl
    Load balance label: Label 17: None;
    Indirect next hop: 0x9e55580 1048574 INH Session ID: 0x14c

    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.10.10.2 via ge-2/0/2.0
        Session Id: 0x146
10.255.162.91/32 Originating RIB: inet.3
    Metric: 1      Node path count: 1
    Forwarding nexthops: 1
    Nexthop: 10.10.10.2 via ge-2/0/2.0

```

user@PE2> **show route table mpls.0**

```

mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:23:33, metric 1
           to table inet.0
0(S=0)     *[MPLS/0] 00:23:33, metric 1
           to table mpls.0
1          *[MPLS/0] 00:23:33, metric 1
           Receive
2          *[MPLS/0] 00:23:33, metric 1
           to table inet6.0
2(S=0)     *[MPLS/0] 00:23:33, metric 1
           to table mpls.0
13         *[MPLS/0] 00:23:33, metric 1
           Receive
17         *[VPN/0] 00:23:33
           to table vpn.inet.0, Pop

```



```

299856(S=0)      *[MPLS/0] 00:23:33
                  to table __1.1.1.1__.mpls.0
299904           *[LDP/9] 00:01:50, metric 1
                  > to 10.10.11.2 via xe-8/2/5.0, Pop
299904(S=0)      *[LDP/9] 00:01:50, metric 1
                  > to 10.10.11.2 via xe-8/2/5.0, Pop
299920           *[LDP/9] 00:01:50, metric 1
                  > to 10.10.11.2 via xe-8/2/5.0, Swap 299904
300016           *[LDP/9] 00:01:50, metric 1
                  > to 10.10.12.1 via ge-3/0/2.0, Pop
                  to table __1.1.1.1__.mpls.0
300016(S=0)      *[LDP/9] 00:01:50, metric 1
                  > to 10.10.12.1 via ge-3/0/2.0, Pop
                  to table __1.1.1.1__.mpls.0
300048           *[LDP/9] 00:01:50, metric 1
                  > to 10.10.12.1 via ge-3/0/2.0, Pop
300048(S=0)      *[LDP/9] 00:01:50, metric 1
                  > to 10.10.12.1 via ge-3/0/2.0, Pop

```

user@PE2> **show route table __1.1.1.1__.mpls.0**

```

__1.1.1.1__.mpls.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

16                *[Egress-Protection/170] 00:22:57
                  to table __1.1.1.1-vpn__.inet.0

```

user@PE2> **show route table __1.1.1.1__.mpls.0 extensive**

```

__1.1.1.1__.mpls.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
16 (1 entry, 1 announced)
    State: < CalcForwarding >
TSI:
KRT in-kernel 16    /52 -> {Table}
    *Egress-Protection Preference: 170
      Next table: __1.1.1.1-vpn__.inet.0
      Next-hop index: 649
      Address: 0x9dc2690

```



```

Next-hop reference count: 2
State: < Active NoReadvrt ForwardingOnly Int Ext >
Local AS: 65000
Age: 22:59
Validation State: unverified
Task: Protection
Announcement bits (1): 0-KRT
AS path: I
Protecting 2 routes

```

user@PE2> **show route table __1.1.1.1-vpn__inet.0**

```

__1.1.1.1-vpn__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.30.0/24      *[Egress-Protection/170] 00:02:11
                   to table vpn.inet.0
10.10.50.0/24      *[Egress-Protection/170] 00:02:11
                   > to 10.10.30.2 via ge-3/2/4.0

```

user@PE2> **show route table __1.1.1.1-vpn__inet.0 extensive**

```

__1.1.1.1-vpn__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.10.30.0/24 (1 entry, 1 announced)
    State: < CalcForwarding >
TSI:
KRT in-kernel 10.10.30.0/24 -> {Table}
    *Egress-Protection Preference: 170
        Next table: vpn.inet.0
        Next-hop index: 592
        Address: 0x9dc2630
        Next-hop reference count: 2
        State: < Active NoReadvrt ForwardingOnly Int Ext >
        Local AS: 65000
        Age: 2:13
        Validation State: unverified
        Task: Protection
        Announcement bits (1): 0-KRT
        AS path: I
        Backup route 10.10.30.0 table vpn.inet.0

```



```

10.10.50.0/24 (1 entry, 1 announced)
    State: < CalcForwarding >
TSI:
KRT in-kernel 10.10.50.0/24 -> {10.10.30.2}
    *Egress-Protection Preference: 170
        Next hop type: Router, Next hop index: 630
        Address: 0x9dc1d90
        Next-hop reference count: 7
        Next hop: 10.10.30.2 via ge-3/2/4.0, selected
        Session Id: 0x147
        State: < Active NoReadvrt ForwardingOnly Int Ext >
        Local AS: 65000
        Age: 2:13
        Validation State: unverified
        Task: Protection
        Announcement bits (1): 0-KRT
        AS path: I
        Backup route 10.10.50.0 table vpn.inet.0

```

user@PE2> **show route table mpls.0 label 17**

```

mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

17                *[VPN/0] 00:25:06
                   to table vpn.inet.0, Pop

```

user@PE2> **show route table mpls.0 label 17 extensive**

```

mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
17 (1 entry, 0 announced)
    *VPN      Preference: 0
        Next table: vpn.inet.0
        Next-hop index: 0
        Label operation: Pop
        Load balance label: None;
        Label element ptr: 0x9db3920
        Label parent element ptr: 0x0
        Label element references: 1
        Label element child references: 0

```



```

Label element lsp id: 0
Address: 0x9db3990
Next-hop reference count: 1
State: < Active NotInstall Int Ext >
Age: 25:30
Validation State: unverified
Task: RT
AS path: I

```

user@PE3> **show route table mpls.0**

```

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:24:16, metric 1
           to table inet.0
0(S=0)     *[MPLS/0] 00:24:16, metric 1
           to table mpls.0
1          *[MPLS/0] 00:24:16, metric 1
           Receive
2          *[MPLS/0] 00:24:16, metric 1
           to table inet6.0
2(S=0)     *[MPLS/0] 00:24:16, metric 1
           to table mpls.0
13         *[MPLS/0] 00:24:16, metric 1
           Receive
16         *[VPN/0] 00:24:15
           to table vpn.inet.0, Pop
300096     *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Swap 299920
300112     *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Swap 299904
300128     *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Pop
300128(S=0) *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Pop

```

user@PE3> **show route table mpls.0 label 16**


```

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

16          *[VPN/0] 00:24:22
           to table vpn.inet.0, Pop

```

user@PE3> **show route table mpls.0 label 16 extensive**

```

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
16  (1 entry, 0 announced)
    *VPN      Preference: 0
              Next table: vpn.inet.0
              Next-hop index: 0
              Label operation: Pop
              Load balance label: None;
              Label element ptr: 0x31d1ec0
              Label parent element ptr: 0x0
              Label element references: 1
              Label element child references: 0
              Label element lsp id: 0
              Address: 0x31d1f30
              Next-hop reference count: 1
              State: < Active NotInstall Int Ext >
              Age: 24:24
              Validation State: unverified
              Task: RT
              AS path: I

```

Checking the Context Identifier Route

Purpose

Examine the information about the context identifier (1.1.1.1).

Action

user@PE1> **show route 1.1.1.1**

```

inet.0: 47 destinations, 47 routes (46 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32      *[IS-IS/15] 00:04:08, metric 31

```



```

> to 10.10.10.2 via ge-2/0/2.0

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[LDP/9] 00:04:08, metric 1
                   > to 10.10.10.2 via ge-2/0/2.0, Push 300064

inet.5: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[IS-IS/15] 00:04:08, metric 31, metric2 1
                   > to 10.10.10.2 via ge-2/0/2.0, Push 299856, Push 299920(top)

```

user@PE2> **show route 1.1.1.1**

```

inet.0: 48 destinations, 49 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[MPLS/2] 00:26:00, metric 16777215
                   Receive
                   [IS-IS/15] 00:04:17, metric 11
                   > to 10.10.12.1 via ge-3/0/2.0

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[LDP/9] 00:04:17, metric 1
                   > to 10.10.12.1 via ge-3/0/2.0

```

user@PE2> **show mpls context-identifier**

ID	Type	Metric	ContextTable
1.1.1.1	protector	16777215	__1.1.1.1__.mpls.0
Total 1, Primary 0, Protector 1			

user@PE2> **show mpls context-identifier detail**

```

ID: 1.1.1.1
  Type: protector, Metric: 16777215, Mode: alias

```



```
Context table: __1.1.1.1__.mpls.0, Label out: 299856
```

```
Total 1, Primary 0, Protector 1
```

user@PE3> **show route 1.1.1.1**

```
inet.0: 47 destinations, 47 routes (46 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.1.1.1/32          *[MPLS/1] 00:26:09, metric 1
                    Receive
```

```
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.1.1.1/32          *[MPLS/1] 00:26:09, metric 1
                    Receive
```

```
inet.5: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.1.1.1/32          *[IS-IS/15] 00:04:27, metric 1, metric2 1
                    > to 10.10.12.2 via ge-1/1/4.0, Push 299856
```

user@PE3> **show mpls context-identifier**

ID	Type	Metric	ContextTable
1.1.1.1	primary	1	
Total 1, Primary 1, Protector 0			

user@PE3> **show mpls context-identifier detail**

```
ID: 1.1.1.1
  Type: primary, Metric: 1, Mode: alias

Total 1, Primary 1, Protector 0
```


Release History Table

Release	Description
15.1	Starting in Junos OS Release 15.1, the enhanced point of local repair (PLR) functionality addresses a special scenario of egress node protection, where the PLR and the protector are co-located as one router. In this case, there is no need to have a bypass LSP reroute traffic during local repair.

RELATED DOCUMENTATION

Egress Protection for Layer 3 VPN Edge Protection Overview

Verifying Path Protection in an MPLS Network

To verify that path protection is working correctly on EX Series switches, perform the following tasks:

1. [Verifying the Primary Path | 234](#)
2. [Verifying the RSVP-Enabled Interfaces | 235](#)
3. [Verifying a Secondary Path | 236](#)

Verifying the Primary Path

Purpose

Verify that the primary path is operational.

Action

```
user@switch> show mpls lsp extensive ingress
```

```
Ingress LSP: 2 sessions
```

```
127.1.8.8
```

```
From: 127.1.9.9, State: Up, ActiveRoute: 0, LSPname: lsp_to_240
```

```
ActivePath: primary_path_lsp_to_240 (primary)
```

```
LoadBalance: Random
```

```
Encoding type: Packet, Switching type: Packet, GPID: IPv4
```

```
*Primary primary_path_lsp_to_240 State: Up
```

```
Priorities: 7 0
```



```

SmartOptimizeTimer: 180
  Exclude: red
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
10.3.3.2 S 10.3.4.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.3.3.2 10.3.4.2
    6 Mar 11 23:58:01.684 Selected as active path: due to 'primary'
    5 Mar 11 23:57:00.750 Record Route:  10.3.3.2 10.3.4.2
    4 Mar 11 23:57:00.750 Up
    3 Mar 11 23:57:00.595 Originate Call
    2 Mar 11 23:57:00.595 CSPF: computation result accepted  10.3.3.2 10.3.4.2
    1 Mar 11 23:56:31.135 CSPF failed: no route toward 10.3.2.2[25 times]
Standby  secondary_path_lsp_to_240 State: Up
Standby  secondary_path_lsp_to_240 State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 1)
10.3.5.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.3.5.2
    7 Mar 11 23:58:01.684 Deselected as active: due to 'primary'
    6 Mar 11 23:46:17.298 Selected as active path
    5 Mar 11 23:46:17.295 Record Route:  5.5.5.2
    4 Mar 11 23:46:17.287 Up
    3 Mar 11 23:46:16.760 Originate Call
    2 Mar 11 23:46:16.760 CSPF: computation result accepted  10.3.5.2
    1 Mar 11 23:45:48.095 CSPF failed: no route toward 10.5.5.5[2 times]
Created: Wed Mar 11 23:44:37 2009
[Output truncated]

```

Meaning

As indicated by the **ActivePath** in the output, the LSP **primary_path_lsp_to_240** is active.

Verifying the RSVP-Enabled Interfaces

Purpose

Verify the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.

Action

```
user@switch> show rsvp interfaces
```



```

RSVP interface: 1 active

```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
ge-0/0/20.0	Up	2	100%	1000Mbps	1000Mbps	0bps	0bps

Meaning

This output verifies that RSVP is enabled and operational on interface **ge-0/0/20.0**.

Verifying a Secondary Path

Purpose

Verify that a secondary path is established.

Action

Deactivate a switch that is critical to the primary path and then issue the following command:

```
user@switch> show mpls lsp extensive
```

```

Ingress LSP: 1 sessions

127.0.0.8
  From: 127.0.0.1, State: Up, ActiveRoute: 0, LSPname: lsp_to_240
  ActivePath: secondary_path_lsp_to_240 (secondary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary   primary_path_lsp_to_240 State: Dn
    Priorities: 7 0
    SmartOptimizeTimer: 180
      Exclude: red
    Will be enqueued for recomputation in 8 second(s).
  51 Mar  8 12:23:31.268 CSPF failed: no route toward 127.0.0.11[11420 times]
  50 Mar  4 15:35:25.610 Clear Call: CSPF computation failed
  49 Mar  4 15:35:25.610 CSPF: link down/deleted:
127.0.0.2(127.0.0.1:0)(127.0.0.1)->
0.0.0.0(127.0.0.20:0)(127.0.0.20)
  48 Mar  4 15:35:25.576 Deselected as active
  47 Mar  4 15:35:25.550 No Route toward dest
  46 Mar  4 15:35:25.550 ??????
  45 Mar  4 15:35:25.549 127.0.0.12: Down
  44 Mar  4 15:33:29.839 Selected as active path

```



```

43 Mar  4 15:33:29.837 Record Route:  127.0.0.20 127.0.0.40
42 Mar  4 15:33:29.835 Up
41 Mar  4 15:33:29.756 Originate Call
40 Mar  4 15:33:29.756 CSPF: computation result accepted  127.0.0.20 127.0.0.40

39 Mar  4 15:33:00.395 CSPF failed: no route toward 127.0.0.11[7 times]
38 Mar  4 15:30:31.412 Clear Call: CSPF computation failed
37 Mar  4 15:30:31.412 CSPF: link down/deleted:
127.0.0.2(127.0.0.1:0)(127.0.0.1)->
0.0.0.0(127.0.0.20:0)(127.0.0.20)
36 Mar  4 15:30:31.379 Deselected as active
35 Mar  4 15:30:31.350 No Route toward dest
34 Mar  4 15:30:31.350 ??????
33 Mar  4 15:30:31.349 127.0.0.12: Down
32 Mar  4 15:29:05.802 Selected as active path
31 Mar  4 15:29:05.801 Record Route:  127.0.0.20 127.0.0.40
30 Mar  4 15:29:05.801 Up
29 Mar  4 15:29:05.686 Originate Call
28 Mar  4 15:29:05.686 CSPF: computation result accepted  127.0.0.20 127.0.0.40

27 Mar  4 15:28:35.852 CSPF failed: no route toward 127.0.0.11[132 times]
26 Mar  4 14:25:12.113 Clear Call: CSPF computation failed
25 Mar  4 14:25:12.113 CSPF: link down/deleted:
0.0.0.0(127.0.0.20:0)(127.0.0.20)->
0.0.0.0(10.10.10.10:0)(10.10.10.10)
*Standby  secondary_path_lsp_to_240 State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 1)
[Output truncated]

```

Meaning

As indicated by the **ActivePath** in the output, the LSP **secondary_path_lsp_to_240** is active.

RELATED DOCUMENTATION

[Configuring Path Protection in an MPLS Network \(CLI Procedure\) | 145](#)

[Understanding MPLS and Path Protection on EX Series Switches | 43](#)

Configuring MPLS Load Balancing and Statistics

IN THIS CHAPTER

- [MPLS Encapsulated Payload Load-balancing Overview | 239](#)
- [Configuring MPLS Encapsulated Payload for Load Balancing | 241](#)
- [Configuring MPLS to Gather Statistics | 242](#)
- [On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)
- [Example: Configuring On-Demand Loss and Delay Measurement | 250](#)
- [Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs | 262](#)
- [Configuring On-Demand Loss and Delay Measurement | 272](#)
- [Configuring Pro-Active Loss and Delay Measurements | 273](#)

MPLS Encapsulated Payload Load-balancing Overview

Routers can load-balance on a per-packet basis in MPLS. Load balancing can be performed on the information in both the IP header and on up to three MPLS labels, providing a more uniform distribution of MPLS traffic to next hops.

Load balancing is used to evenly distribute traffic when the following conditions apply:

- There are multiple equal-cost next hops over different interfaces to the same destination.
- There is a single next hop over an aggregated interface.

By default, when load balancing is used to help distribute traffic, a hash algorithm is used to select a next-hop address to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is reselected by means of the hash algorithm.

In case of multiple transport layer networks such as Ethernet over MPLS or Ethernet pseudowire, the hash algorithm needs to look beyond the outer header of the payload and into the inner headers to generate an even distribution. To determine the inner encapsulation, the PFE relies on the presence of certain codes or numbers at fixed payload offsets; for example the presence of payload type 0X800 or the presence of protocol number 4 for an IPv4 packet. In Junos OS, you can configure **zero-control-word** option to indicate the start of an Ethernet frame in an MPLS ether-pseudowire payload. On seeing this control word, which

is four bytes having a numerical value of all zeros, the hash generator assumes the start of an Ethernet frame at the end of the control word in an MPLS ether-pseudowire packet.

NOTE: For DPC I-chip-based cards, configure the **zero-control-word** option at the **[edit forwarding-options hash-key family mpls ether-pseudowire]** hierarchy level; and for MPC cards, configure the **zero-control-word** option at the **[edit forwarding-options enhanced-hash-key family mpls ether-pseudowire]** hierarchy level.

RELATED DOCUMENTATION

Configuring MPLS Encapsulated Payload for Load Balancing 241
enhanced-hash-key
ether-pseudowire 2605
hash-key
family mpls 2615

Configuring MPLS Encapsulated Payload for Load Balancing

By default, when load balancing is used to help distribute traffic, a hash algorithm is used to select a next-hop address to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is reselected by means of the hash algorithm. Configure the **zero-control-word** option to indicate the start of an Ethernet frame in an MPLS ether-pseudowire payload. On seeing this control word, four bytes having a numerical value of all zeros, the hash generator assumes the start of the Ethernet frame at the end of the control word in an MPLS ether-pseudowire packet.

Before you begin to configure MPLS encapsulated payload for load balancing, configure routing and signaling protocols.

To configure MPLS encapsulated payload for load balancing:

1. Configure the **zero-control-word** option to indicate the start of an Ethernet frame in an MPLS ether-pseudowire payload.
 - For DPC I-chip-based cards, configure the **zero-control-word** option at the [edit forwarding-options hash-key family mpls ether-pseudowire] hierarchy level.

```
[edit forwarding-options hash-key family mpls ether-pseudowire]
user@host# set zero-control-word
```

- For MPC cards, configure the **zero-control-word** option at the [edit forwarding-options enhanced-hash-key family mpls ether-pseudowire] hierarchy level.

```
[edit forwarding-options enhanced-hash-key family mpls ether-pseudowire]
user@host# set zero-control-word
```

RELATED DOCUMENTATION

MPLS Encapsulated Payload Load-balancing Overview 239
enhanced-hash-key
ether-pseudowire 2605
hash-key
family mpls 2615

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 (MPLS Statistics);
  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 | 547](#)

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

IN THIS SECTION

- [Importance of Measuring Packet Loss and Delay | 243](#)
- [Defining Packet Loss, Delay, and Throughput | 244](#)
- [Packet Loss and Delay Measurement Mechanisms | 244](#)
- [Packet Loss and Delay Metrics | 245](#)
- [Packet Loss and Delay Measurement Concepts | 245](#)
- [Packet Loss and Delay Measurement Functionality | 248](#)
- [Packet Loss and Delay Features | 249](#)

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

Importance of Measuring Packet Loss and Delay

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

With the emergence of new applications on data networks, it is becoming increasingly important for service providers to accurately predict the impact of new application rollouts. Understanding and modelling network performance in the network is especially relevant for deployment of new-world applications to ensure

successful implementations. In packet networks, packet loss and delay are two of the most fundamental measures of performance. Their role is even more central when it comes to end-to-end measurements.

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

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

Defining Packet Loss, Delay, and Throughput

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

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

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

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

A low latency indicates a high network efficiency.

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

Packet Loss and Delay Measurement Mechanisms

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

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

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

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

Packet Loss and Delay Metrics

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

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

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

Packet Loss and Delay Measurement Concepts

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

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

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

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

- **Measurement point (MP)**—MP is the location at which a condition is described for the measurement.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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


- **Timestamp format**—Junos OS supports only the IEEE 1588 Precision Time Protocol (PTP) [IEEE1588] format for recording delay measurement messages. Network Time Format (NTP) is not supported.
- **Operations, administration, and maintenance (OAM)**—To indicate that all the OAM messages for MPLS LSPs flow over the MPLS G-Ach, and to enable the MPLS performance monitoring messages to be carried

over the MPLS G-Ach, the **oam mpls-tp-mode** statement must be included at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level.

Packet Loss and Delay Measurement Functionality

Figure 14 on page 248 illustrates the basic methods used for the bidirectional measurement of packet loss and delay. A bidirectional channel exists between the two routers, Router A and Router B. The temporal reference points – T1, T2, T3, and T4 – are associated with a measurement operation that takes place at Router A. The operation consists of Router A sending a query message to Router B, and Router B sending back a response. Each reference point indicates the point of time at which either the query or the response message is transmitted or received over the channel.

Figure 14: Basic Bidirectional Measurement



In Figure 14 on page 248, Router A can arrange to measure the packet loss over the channel in the forward and reverse directions by sending loss measurement query messages to Router B. Each of the forward and reverse messages contain the count of packets transmitted prior to time T1 over the channel to Router B (A_TxP).

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

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

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

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

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

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

The arithmetic is modulo the counter size.

To measure at Router A the delay over the channel to Router B, a delay measurement query message is sent from Router A to Router B containing a timestamp recording the instant at which it is transmitted. In [Figure 14 on page 248](#), the timestamp is recorded in T1.

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

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

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

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

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

Packet Loss and Delay Features

Supported Features of Packet Loss and Delay

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

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

Unsupported Features of Packet Loss and Delay

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

- Loss and delay measurement for pseudowires (section 2.9.1 of RFC 6374)
- Unidirectional measurement (section 2.6 of RFC 6374)
- Dyadic measurement (section 2.7 of RFC 6374)

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

RELATED DOCUMENTATION

[Example: Configuring On-Demand Loss and Delay Measurement | 250](#)

[monitor mpls loss rsvp | 3114](#)

[monitor mpls delay rsvp | 3108](#)

[monitor mpls loss-delay rsvp | 3120](#)

Example: Configuring On-Demand Loss and Delay Measurement

IN THIS SECTION

- [Requirements | 250](#)
- [Overview | 251](#)
- [Configuration | 251](#)
- [Verification | 256](#)

This example shows how to enable on-demand loss and delay measurement for point-to-point ultimate hop popping (UHP) label-switched paths (LSPs) in MPLS networks to monitor network performance.

Requirements

This example uses the following hardware and software components:

- Two MX Series 5G Universal Routing Platforms that contain MPC/MICs only

- Junos OS Release 14.2 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
 - RSVP
 - MPLS
 - OSPF

Overview

Starting with Junos OS Release 14.2, an on-demand tool to monitor and measure packet loss, packet delay, or both for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs) is introduced. The tool can be enabled using the following CLI commands – **monitor mpls loss rsvp**, **monitor mpls delay rsvp**, and **monitor mpls loss-delay rsvp**.


These commands provide an on-demand summary of performance metrics for direct mode packet loss, two-way packet delay, and related metrics, such as inter-packet delay variation and channel throughput measurement.

This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

Topology

Figure 15 on page 251 illustrates the on-demand loss and delay measurement using a simple two-router topology.

Figure 15: Configuring On-Demand Loss and Delay Measurement



In this example, an associated bidirectional LSP is configured between Routers R1 and R2, for which the performance metrics is measured.

Configuration

CLI Quick Configuration

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

R1

```
set chassis fpc 0 pic 3 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.0.0.1
set protocols rsvp interface ge-0/0/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls statistics traffic-class-statistics
set protocols mpls label-switched-path R1-R2 to 20.0.0.1
set protocols mpls label-switched-path R1-R2 oam mpls-tp-mode
set protocols mpls label-switched-path R1-R2 ultimate-hop-popping
set protocols mpls label-switched-path R1-R2 associate-lsp R2-R1
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

R2

```
set chassis fpc 0 pic 3 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 20.0.0.1/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 20.0.0.1
set protocols rsvp interface ge-0/0/0.0
set protocols rsvp interface lo0.0
```



```

set protocols rsvp interface fxp0.0 disable
set protocols mpls statistics traffic-class-statistics
set protocols mpls label-switched-path R2-R1 to 10.0.0.1
set protocols mpls label-switched-path R2-R1 oam mpls-tp-mode
set protocols mpls label-switched-path R2-R1 ultimate-hop-popping
set protocols mpls label-switched-path R2-R1 associate-lsp R1-R2
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0.interface fxp0.0 disable

```

Step-by-Step Procedure

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

To configure Router R1:

1. Enable the chassis with tunnel services and enhanced IP network services configuration.

```

[edit chassis]
user@R1# set fpc 0 pic 3 tunnel-services bandwidth 1g
user@R1# set network-services enhanced-ip

```

2. Configure the interfaces for Router R1.

```

[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet address 1.1.1.1/30
user@R1# set ge-0/0/0 unit 0 family mpls
user@R1# set lo0 unit 0 family inet address 10.0.0.1/32
user@R1# set lo0 unit 0 family mpls

```

3. Configure the router ID for Router R1.

```

[edit routing-options]
user@R1# set router-id 10.0.0.1

```


4. Enable RSVP on all the interfaces of Router R1, excluding the management interface.

```
[edit protocols]
user@R1# set rsvp interface ge-0/0/0.0
user@R1# set rsvp interface lo0.0
user@R1# set rsvp interface fxp0.0 disable
```

5. Enable MPLS on all the interfaces of Router R1, excluding the management interface.

```
[edit protocols]
user@R1# set mpls interface ge-0/0/0.0
user@R1# set mpls interface lo0.0
user@R1# set mpls interface fxp0.0 disable
```

6. Configure an associated bidirectional LSP to Router R2.

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 to 20.0.0.1
user@R1# set mpls label-switched-path R1-R2 oam mpls-tp-mode
user@R1# set mpls label-switched-path R1-R2 ultimate-hop-popping
user@R1# set mpls label-switched-path R1-R2 associate-lsp R2-R1
```

7. Create traffic classes for maintaining data traffic statistics per traffic class.

This enables traffic class scoped loss measurement.

```
[edit protocols]
user@R1# set mpls statistics traffic-class-statistics
```

8. Configure OSPF with traffic engineering capabilities, and enable OSPF on all the interfaces of Router R1, excluding the management interface.

```
[edit protocols]
user@R1# set ospf traffic-engineering
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@R1# set ospf area 0.0.0.0 interface lo0.0
user@R1# set ospf interface fxp0.0 disable
```


Results

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

```
user@R1# show chassis
fpc 0 {
  pic 3 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}
network-services enhanced-ip;
```

```
user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 1.1.1.1/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.0.0.1/32;
    }
    family mpls;
  }
}
```

```
user@R1# show routing-options
router-id 10.0.0.1;
```

```
user@R1# show protocols
rsvp {
  interface ge-0/0/0.0;
  interface lo0.0;
  interface fxp0.0 {
```



```

        disable;
    }
}
mpls {
    statistics {
        traffic-class-statistics;
    }
    label-switched-path R1-R2 {
        to 20.0.0.1;
        oam mpls-tp-mode;
        ultimate-hop-popping;
        associate-lsp R2-R1;
    }
    interface ge-0/0/0.0;
    interface lo0.0;
    interface fxp0.0 {
        disable;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/0/0.0;
        interface lo0.0;
        interface fxp0.0 {
            disable;
        }
    }
}
}

```

Verification

IN THIS SECTION

- [Verifying the LSP Status | 257](#)
- [Verifying Packet Loss Measurement | 257](#)
- [Verifying Packet Delay Measurement | 259](#)
- [Verifying Packet Loss-Delay Measurement | 260](#)

Confirm that the configuration is working properly.

Verifying the LSP Status

Purpose

Verify that the associated bidirectional LSP between Routers R1 and R2 is up.

Action

From operational mode, run the **show mpls lsp** command.

```
user@R1> show mpls lsp
```

```
Ingress LSP: 1 sessions
To          From          State Rt P    ActivePath      LSPname
20.0.0.1    10.0.0.1    Up     0  *              R1-R2 Assoc-Bidir
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPname
10.0.0.1    20.0.0.1    Up     0  1 FF  299776      - R2-R1 Assoc-Bidir
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

The associated bidirectional LSP R1-R2 is up and active.

Verifying Packet Loss Measurement

Purpose

Verify the on-demand loss measurement result.

Action

From operational mode, run the **monitor mpls loss rsvp R1-R2 count 2 detail** command.

```
user@R1> monitor mpls loss rsvp R1-R2 count 2 detail
```

```
(0)
Response code                : Success
Origin timestamp              : 1404129082 secs, 905571890 nsecs
Forward transmit count        : 83040
Forward receive count         : 83040
```



```

Reverse transmit count      : 83100
Reverse receive count      : 83100
(1)
Response code              : Success
Origin timestamp           : 1404129083 secs, 905048410 nsecs
Forward transmit count     : 83841
Forward receive count      : 83841
Reverse transmit count     : 83904
Reverse receive count      : 83904
Current forward transmit count : 801
Current forward receive count : 801
Current forward loss       : 0 packets
Current forward loss ratio  : 0.000000
Current forward throughput  : 0.801 kpps
Current reverse transmit count : 804
Current reverse receive count : 804
Current reverse loss       : 0 packets
Current reverse loss ratio  : 0.000000
Current reverse throughput  : 0.804 kpps
(2)
Response code              : Success
Origin timestamp           : 1404129084 secs, 904828715 nsecs
Forward transmit count     : 84423
Forward receive count      : 84423
Reverse transmit count     : 84487
Reverse receive count      : 84487
Current forward transmit count : 582
Current forward receive count : 582
Current forward loss       : 0 packets
Current forward loss ratio  : 0.000000
Current forward throughput  : 0.582 kpps
Current reverse transmit count : 583
Current reverse receive count : 583
Current reverse loss       : 0 packets
Current reverse loss ratio  : 0.000000
Current reverse throughput  : 0.583 kpps

Cumulative forward transmit count : 1383
Cumulative forward loss       : 0 packets
Average forward loss ratio    : 0.000000
Average forward throughput    : 0.692 kpps
Cumulative reverse transmit count : 1387
Cumulative reverse loss       : 0 packets
Average reverse loss ratio    : 0.000000

```



```

Average reverse throughput      : 0.694 kpps

LM queries sent                : 3
LM responses received          : 3
LM queries timedout            : 0
LM responses dropped due to errors : 0

```

Meaning

The packet loss measurement for two counts is displayed.

Verifying Packet Delay Measurement

Purpose

Verify the on-demand delay measurement result.

Action

From operational mode, run the **monitor mpls delay rsvp R1-R2 count 2 detail** command.

```
user@R1> monitor mpls delay rsvp R1-R2 count 2 detail
```

```

(1)
Response code                : Success
Querier transmit timestamp   : 1404129122 secs, 479955401 nsecs
Responder receive timestamp  : 1404129122 secs, 468519022 nsecs
Responder transmit timestamp : 1404129122 secs, 470255123 nsecs
Querier receive timestamp    : 1404129122 secs, 481736403 nsecs
Current two-way channel delay : 44 usecs
Current round-trip-time      : 1781 usecs

(2)
Response code                : Success
Querier transmit timestamp   : 1404129123 secs, 480926210 nsecs
Responder receive timestamp  : 1404129123 secs, 469488696 nsecs
Responder transmit timestamp : 1404129123 secs, 471130706 nsecs
Querier receive timestamp    : 1404129123 secs, 482613911 nsecs
Current two-way channel delay : 45 usecs
Current round-trip-time      : 1687 usecs

Best two-way channel delay   : 44 usecs
Worst two-way channel delay  : 45 usecs
Average two-way channel delay : 45 usecs
Best round-trip-time         : 1687 usecs
Worst round-trip-time        : 1781 usecs
Average round-trip-time      : 1734 usecs

```



```

Average forward delay variation      : 1 usecs
Average reverse delay variation      : 1 usecs

DM queries sent                     : 2
DM responses received                : 2
DM queries timedout                  : 0
DM responses dropped due to errors   : 0

```

Meaning

The packet delay measurement for two counts is displayed.

Verifying Packet Loss-Delay Measurement

Purpose

Verify the on-demand loss and delay measurement result.

Action

From operational mode, run the **monitor mpls loss-delay rsvp R1-R2 count 2 detail** command.

```
user@R1> monitor mpls loss-delay rsvp R1-R2 count 2 detail
```

```

(0)
Response code                      : Success
Forward transmit count              : 142049
Forward receive count               : 142049
Reverse transmit count              : 142167
Reverse receive count               : 142167
Querier transmit timestamp          : 1404129161 secs, 554422723 nsecs
Responder receive timestamp         : 1404129161 secs, 542877570 nsecs
Responder transmit timestamp        : 1404129161 secs, 546004545 nsecs
Querier receive timestamp           : 1404129161 secs, 557599327 nsecs
(1)
Response code                      : Success
Forward transmit count              : 143049
Forward receive count               : 143049
Reverse transmit count              : 143168
Reverse receive count               : 143168
Current forward transmit count      : 1000
Current forward receive count       : 1000
Current forward loss                 : 0 packets
Current forward loss ratio          : 0.000000
Current forward throughput          : 1.000 kpps
Current reverse transmit count      : 1001

```



```

Current reverse receive count      : 1001
Current reverse loss               : 0 packets
Current reverse loss ratio        : 0.000000
Current reverse throughput        : 1.001 kpps
Querier transmit timestamp        : 1404129162 secs, 554465742 nsecs
Responder receive timestamp       : 1404129162 secs, 542919166 nsecs
Responder transmit timestamp      : 1404129162 secs, 545812736 nsecs
Querier receive timestamp         : 1404129162 secs, 557409175 nsecs
Current two-way channel delay     : 49 usecs
Current round-trip-time           : 2943 usecs
(2)
Response code                     : Success
Forward transmit count            : 143677
Forward receive count             : 143677
Reverse transmit count            : 143799
Reverse receive count             : 143799
Current forward transmit count    : 628
Current forward receive count     : 628
Current forward loss              : 0 packets
Current forward loss ratio       : 0.000000
Current forward throughput       : 0.627 kpps
Current reverse transmit count    : 631
Current reverse receive count     : 631
Current reverse loss              : 0 packets
Current reverse loss ratio       : 0.000000
Current reverse throughput       : 0.630 kpps
Querier transmit timestamp        : 1404129163 secs, 556698575 nsecs
Responder receive timestamp       : 1404129163 secs, 545150128 nsecs
Responder transmit timestamp      : 1404129163 secs, 546918408 nsecs
Querier receive timestamp         : 1404129163 secs, 558515047 nsecs
Current two-way channel delay     : 48 usecs
Current round-trip-time           : 1816 usecs

Cumulative forward transmit count : 1628
Cumulative forward loss           : 0 packets
Average forward loss ratio        : 0.000000
Average forward throughput        : 0.813 kpps
Cumulative reverse transmit count : 1632
Cumulative reverse loss           : 0 packets
Average reverse loss ratio        : 0.000000
Average reverse throughput        : 0.815 kpps

Best two-way channel delay        : 48 usecs
Worst two-way channel delay       : 49 usecs

```


Average two-way channel delay	: 49 usecs
Best round-trip-time	: 1816 usecs
Worst round-trip-time	: 3176 usecs
Average round-trip-time	: 2645 usecs
Average forward delay variation	: 1 usecs
Average reverse delay variation	: 0 usecs
LDM queries sent	: 3
LDM responses received	: 3
LDM queries timedout	: 0
LDM responses dropped due to errors	: 0

Meaning

The packet loss and delay measurement for two counts is displayed.

RELATED DOCUMENTATION

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview 243
monitor mpls loss rsvp 3114
monitor mpls delay rsvp 3108
monitor mpls loss-delay rsvp 3120

Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs

IN THIS SECTION

- [Requirements | 263](#)
- [Overview | 263](#)
- [Configuration | 264](#)
- [Verification | 270](#)

This example shows how to configure pro-active loss and delay measurements for point-to-point ultimate-hop popping label-switched paths (LSPs) in MPLS networks to monitor network performance.

Requirements

This example uses the following hardware and software components:

- Two MX Series 5G Universal Routing Platforms that contain MPC/MICs only
- Junos OS Release 15.1 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
 - a. MPLS
 - b. OSPF
 - c. RSVP

Overview

Starting with Junos OS Release 15.1, a pro-active tool to monitor and measure packet loss, packet delay, or both for associated bidirectional MPLS ultimate-hop popping point-to-point label-switched paths (LSPs) is introduced.

This feature provides the following performance metrics:


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

This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

Topology

[Figure 15 on page 251](#) illustrates the pro-active loss and delay measurements using a simple two-router topology.

Figure 16: Configuring Pro-Active Loss and Delay Measurements



In this example, an associated bidirectional LSP is configured between Routers R1 and R2, for which the performance metrics are measured.

Configuration

CLI Quick Configuration

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

R1

```

set chassis network-services enhanced-ip
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set interfaces lo0 unit 0 family mpls
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols mpls label-switched-path R1-R2 associate-lsp R2-R1
set protocols mpls label-switched-path R1-R2 install 20.10.30.0/24 active
set protocols mpls label-switched-path R1-R2 oam mpls-tp-mode
set protocols mpls label-switched-path R1-R2 oam performance-monitoring querier delay traffic-class
  tc-0 query-interval 1000
set protocols mpls label-switched-path R1-R2 oam performance-monitoring querier loss traffic-class
  none query-interval 1000
set protocols mpls label-switched-path R1-R2 oam performance-monitoring querier loss-delay
  traffic-class tc-0 query-interval 1000
set protocols mpls label-switched-path R1-R2 oam performance-monitoring responder delay
  min-query-interval 1000
set protocols mpls label-switched-path R1-R2 oam performance-monitoring responder loss
  min-query-interval 1000
set protocols mpls label-switched-path R1-R2 to 20.0.0.1
set protocols mpls label-switched-path R1-R2 ultimate-hop-popping
set protocols mpls statistics traffic-class-statistics
  
```



```

set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols rsvp interface ge-0/0/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set routing-options router-id 10.0.0.1

```

R2

```

set chassis network-services enhanced-ip
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 20.0.0.1/32
set interfaces lo0 unit 0 family mpls
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols mpls label-switched-path R2-R1 associate-lsp R1-R2
set protocols mpls label-switched-path R2-R1 install 10.10.20.0/24 active
set protocols mpls label-switched-path R2-R1 oam mpls-tp-mode
set protocols mpls label-switched-path R2-R1 oam performance-monitoring responder delay
    min-query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring responder loss
    min-query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring querier delay traffic-class
    tc-0 query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring querier loss traffic-class
    none query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring querier loss-delay
    traffic-class tc-0 query-interval 1000
set protocols mpls label-switched-path R2-R1 to 10.0.0.1
set protocols mpls label-switched-path R2-R1 ultimate-hop-popping
set protocols mpls statistics traffic-class-statistics
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols rsvp interface ge-0/0/0.0

```



```

set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set routing-options router-id 20.0.0.1

```

Step-by-Step Procedure

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

To configure Router R1:

1. Enable the enhanced IP network services configuration.

```

[edit chassis]
user@R1# set network-services enhanced-ip

```

2. Configure the interfaces for Router R1.

```

[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet address 1.1.1.1/30
user@R1# set ge-0/0/0 unit 0 family mpls
user@R1# set lo0 unit 0 family inet address 10.0.0.1/32
user@R1# set lo0 unit 0 family mpls

```

3. Configure the router ID for Router R1.

```

[edit routing-options]
user@R1# set router-id 10.0.0.1

```

4. Enable RSVP on all the interfaces of Router R1, excluding the management interface.

```

[edit protocols]
user@R1# set rsvp interface ge-0/0/0.0
user@R1# set rsvp interface lo0.0
user@R1# set rsvp interface fxp0.0 disable

```

5. Enable MPLS on all the interfaces of Router R1, excluding the management interface.


```
[edit protocols]
user@R1# set mpls interface ge-0/0/0.0
user@R1# set mpls interface lo0.0
user@R1# set mpls interface fxp0.0 disable
```

6. Configure an associated bidirectional LSP to Router R2.

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 to 20.0.0.1
user@R1# set mpls label-switched-path R1-R2 install 20.10.30.0/24 active
user@R1# set mpls label-switched-path R1-R2 oam mpls-tp-mode
user@R1# set mpls label-switched-path R1-R2 ultimate-hop-popping
user@R1# set mpls label-switched-path R1-R2 associate-lsp R2-R1
```

7. Create traffic classes for maintaining data traffic statistics per traffic class.

This enables traffic class scoped loss and delay measurement.

```
[edit protocols]
user@R1# set mpls statistics traffic-class-statistics
```

8. Configure performance monitoring at the querier side.

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring querier delay traffic-class tc-0
    query-interval 1000
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring querier loss traffic-class none
    query-interval 1000
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring querier loss-delay traffic-class
    tc-0 query-interval 1000
```

9. Configure performance monitoring at the responder side.

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring responder delay
    min-query-interval 1000
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring responder loss min-query-interval
    1000
```


10. Configure OSPF with traffic engineering capabilities, and enable OSPF on all the interfaces of Router R1, excluding the management interface.

```
[edit protocols]
user@R1# set ospf traffic-engineering
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@R1# set ospf area 0.0.0.0 interface lo0.0
user@R1# set ospf interface fxp0.0 disable
```

Results

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

```
user@R1# show chassis
network-services enhanced-ip;
```

```
user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 1.1.1.1/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.0.0.1/32;
    }
    family mpls;
  }
}
```

```
user@R1# show routing-options
router-id 10.0.0.1;
```

```
user@R1# show protocols
```



```

rsvp {
  interface ge-0/0/0.0;
  interface lo0.0;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  label-switched-path R1-R2 {
    to 20.0.0.1;
    install 20.10.30.0/24 active;
    oam {
      mpls-tp-mode;
      performance-monitoring {
        querier {
          loss {
            traffic-class none {
              query-interval 1000;
            }
          }
          delay {
            traffic-class tc-0 {
              query-interval 1000;
            }
          }
          loss-delay {
            traffic-class none {
              query-interval 1000;
            }
          }
        }
      }
      responder {
        loss {
          min-query-interval 1000;
        }
        delay {
          min-query-interval 1000;
        }
      }
    }
  }
  ultimate-hop-popping;
  associate-lsp R2-R1;
}

```



```

}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
    interface fxp0.0 {
      disable;
    }
  }
}
}

```

Verification

Verifying Loss and Delay Measurement

Purpose

Verify the loss and delay measurement.

Action

From operational mode, run the **show performance-monitoring mpls lsp** command.

```
user@R1> show performance-monitoring mpls lsp
```

```

Session Total: 3 Up: 3 Down: 0
LSP name:R1-R2, PM State:Up
Loss measurement Data:
  Duration: 00:04:43
  Traffic-class: None
  Queries sent: 282
  Responses received: 282
  Responses dropped due to errors: 0
  Queries timeout: 0
  Forward loss measurement:
    Average packet loss: 0
    Average packet throughput: 554338
  Reverse loss measurement:
    Average packet loss: 0
    Average packet throughput: 1352077
LSP name:R1-R2, PM State:Up
Delay measurement Data:
  Duration: 00:04:43
  Traffic-class: 0

```



```

Queries sent: 282
Responses received: 282
Responses dropped due to errors: 0
Queries timeout: 0
Best 2-way channel delay: 72 usecs
Worst 2-way channel delay: 365 usecs
Best round trip time: 843 usecs
Worst round trip time: 105523 usecs
Avg absolute fw delay variation: 1619 usecs
Avg absolute rv delay variation: 1619 usecs
LSP name:R1-R2, PM State:Up
Loss measurement Data:
  Duration: 00:04:43
  Traffic-class: None
  Queries sent: 282
  Responses received: 282
  Responses dropped due to errors: 0
  Queries timeout: 0
  Forward loss measurement:
    Average packet loss: 0
    Average packet throughput: 553927
  Reverse loss measurement:
    Average packet loss: 0
    Average packet throughput: 1351531
Delay measurement Data:
  Best 2-way channel delay: 76 usecs
  Worst 2-way channel delay: 368 usecs
  Best round trip time: 1082 usecs
  Worst round trip time: 126146 usecs
  Avg absolute fw delay variation: 1618 usecs
  Avg absolute rv delay variation: 1619 usecs

```

Meaning

The packet loss and delay measurement metrics for LSP are displayed.

RELATED DOCUMENTATION

[performance-monitoring \(Protocols MPLS\) | 2727](#)

[show performance-monitoring mpls lsp | 3254](#)

Configuring On-Demand Loss and Delay Measurement

You can configure an on-demand loss and delay measurement for point-to-point ultimate hop popping (UHP) label-switched paths (LSPs) in MPLS networks to monitor network performance. The **monitor mpls loss rsvp**, **monitor mpls delay rsvp**, and **monitor mpls loss-delay rsvp** CLI commands provide an on-demand summary of performance metrics for direct mode packet loss, two-way packet delay, and related metrics, such as inter-packet delay variation and channel throughput measurement.

This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

Before you begin:

1. Configure the device interfaces.
2. Configure the device router ID.
3. Configure the following protocols:
 - RSVP
 - OSPF

Enable traffic engineering capabilities.

 - MPLS

To configure the PE device:

1. Enable the chassis with tunnel services and enhanced IP network services configuration.

```
[edit chassis]
user@R1# set fpc fpc-slot pic pic-slot tunnel-services bandwidth bandwidth
user@R1# set network-services enhanced-ip
```

2. Configure an associated bidirectional LSP to the remote router.

```
[edit protocols]
user@R1# set mpls label-switched-path lsp-name to remote-router-ip-address
user@R1# set mpls label-switched-path lsp-name oam mpls-tp-mode
user@R1# set mpls label-switched-path lsp-name ultimate-hop-popping
user@R1# set mpls label-switched-path lsp-name associate-lsp lsp-name
```

3. Create traffic classes for maintaining data traffic statistics per traffic class.

This enables traffic class scoped loss measurement.


```
[edit protocols]
user@R1# set mpls statistics traffic-class-statistics
```

RELATED DOCUMENTATION

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[Example: Configuring On-Demand Loss and Delay Measurement | 250](#)

[monitor mpls loss rsvp | 3114](#)

[monitor mpls delay rsvp | 3108](#)

[monitor mpls loss-delay rsvp | 3120](#)

Configuring Pro-Active Loss and Delay Measurements

You can configure pro-active loss and delay measurements for point-to-point ultimate-hop popping label-switched paths (LSPs) in MPLS networks to monitor network performance. The **show performance-monitoring mpls lsp** CLI command provides a summary of performance metrics for direct mode packet loss, two-way packet delay, and related metrics, such as inter-packet delay variation and channel throughput measurement.

This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

This feature provides the following performance metrics:

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

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
 - MPLS

- OSPF
- RSVP

To configure pro-active loss and delay measurements on the PE device:

1. Configure an associated bidirectional LSP to Router R2.

```
[edit protocols]
user@host# set mpls label-switched-path lsp-name to remote-router-ip-address
user@host# set mpls label-switched-path lsp-name install destination-prefix/prefix-length active
user@host# set mpls label-switched-path lsp-name oam mpls-tp-mode
user@host# set mpls label-switched-path lsp-name ultimate-hop-popping
user@host# set mpls label-switched-path lsp-name associate-lsp remote-lsp-name
```

2. Create traffic classes for maintaining data traffic statistics per traffic class.

This enables traffic class scoped loss and delay measurements.

```
[edit protocols]
user@host# set mpls statistics traffic-class-statistics
```

3. Configure performance monitoring at the querier side.

```
[edit protocols]
user@host# set mpls label-switched-path lsp-name oam performance-monitoring querier delay traffic-class
    tc-value query-interval milliseconds
user@host# set mpls label-switched-path lsp-name oam performance-monitoring querier loss traffic-class
    tc-value query-interval milliseconds
user@host# set mpls label-switched-path lsp-name oam performance-monitoring querier loss-delay traffic-class
    tc-value query-interval milliseconds
```

4. Configure performance monitoring at the responder side.

```
[edit protocols]
user@host# set mpls label-switched-path lsp-name oam performance-monitoring responder delay
    min-query-interval milliseconds
user@host# set mpls label-switched-path lsp-name oam performance-monitoring responder loss
    min-query-interval milliseconds
```

RELATED DOCUMENTATION

Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs | 262

performance-monitoring (Protocols MPLS) | 2727

show performance-monitoring mpls lsp | 3254

Configuring Shared Risk Link Group (SRLG)

IN THIS CHAPTER

- [SRLG Overview | 277](#)
- [Example: Configuring SRLG | 278](#)
- [Example: Excluding SRLG Links Completely for the Secondary LSP | 290](#)
- [Example: Configuring SRLG with Link Protection | 298](#)
- [Example: Configuring SRLG with Link Protection with the exclude-srlg Option | 325](#)

SRLG Overview

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

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

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

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

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

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

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

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

RELATED DOCUMENTATION

[Example: Configuring SRLG | 278](#)

[Example: Excluding SRLG Links Completely for the Secondary LSP | 290](#)

[Example: Configuring SRLG with Link Protection | 298](#)

[Example: Configuring SRLG with Link Protection with the exclude-srlg Option | 325](#)

[Computing Backup Paths for LSPs Using Fate Sharing | 504](#)

Example: Configuring SRLG

IN THIS SECTION

- [Requirements | 278](#)
- [Overview | 279](#)
- [Configuration | 280](#)
- [Verification | 287](#)

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

Requirements

This example uses the following hardware and software components:

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

Configuration

CLI Quick Configuration

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

Router PE1

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

Router P1

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.2/24
set interfaces ge-0/0/1 unit 0 family mpls
```



```

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

```

Router P2

```

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

```

Router P3

```

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

```



```

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

```

Router P4

```

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

```

Router P5

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.56.6/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.67.6/24
set interfaces ge-0/0/2 unit 0 family mpls

```



```

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

```

Router PE2

```

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

```

Step-by-Step Procedure

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

To configure the ingress router PE1:

1. Configure the device interfaces.

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

2. Configure OSPF on the interfaces.

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

3. Configure the SRLG definitions.

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

4. Configure MPLS and the LSPs.

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



```
user@PE1# set path path2
```

5. Enable RSVP on the interfaces.

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

Results

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

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



```

        family inet {
            address 10.255.0.1/32;
        }
    }
}

```

```

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

```

```

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

```

```

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

```

```

user@PE1# show routing-options
routing-options {
    srlg {

```



```
    srlg-a {  
      srlg-value 101;  
      srlg-cost 10;  
    }  
  }  
}
```

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

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

Verification

IN THIS SECTION

- [Verifying SRLG Definitions | 287](#)
- [Verify TE-Link SRLG | 288](#)
- [Verify Standby Secondary Path | 288](#)

Confirm that the configuration is working properly.

Verifying SRLG Definitions

Purpose

Verify SRLG-to-value mappings and SRLG cost.

Action

user@PE1> **show mpls srlg**

SRLG	Value	Cost
srlg-a	101	10

Verify TE-Link SRLG**Purpose**

Verify the traffic engineering link SRLG association.

Action

```
user@PE1> show ted link detail
```

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

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

...
```

Meaning

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

Verify Standby Secondary Path**Purpose**

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

Action

```
user@PE1> show mpls lsp ingress extensive
```

```
Ingress LSP: 1 sessions
```



```

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

```



```

3 Oct 13 15:15:14.984 Up
2 Oct 13 15:15:14.830 Originate Call
1 Oct 13 15:15:14.830 CSPF: computation result accepted 192.168.13.3
192.168.37.7
Created: Thu Oct 13 15:13:46 2011
Total 1 displayed, Up 1, Down 0

```

Meaning

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

RELATED DOCUMENTATION

[SRLG Overview | 277](#)

[Example: Excluding SRLG Links Completely for the Secondary LSP | 290](#)

[srlg | 2792](#)

[srlg-cost | 2793](#)

[srlg-value | 2794](#)

Example: Excluding SRLG Links Completely for the Secondary LSP

IN THIS SECTION

- [Requirements | 291](#)
- [Overview | 291](#)
- [Configuration | 292](#)
- [Verification | 296](#)

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

Requirements

This example uses the following hardware and software components:

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

Overview

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

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

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


```

set protocols mpls optimize-timer 120
set protocols mpls exclude-srlg
set protocols mpls label-switched-path pe1-pe2 to 10.255.0.7
set protocols mpls label-switched-path pe1-pe2 primary via-p1
set protocols mpls label-switched-path pe1-pe2 secondary path2 standby
set protocols mpls path via-p1 10.255.0.2 strict
set protocols mpls path path2
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Step-by-Step Procedure

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

1. Configure the device interfaces.

```

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

```

2. Configure OSPF on the interfaces.

```

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

```


3. Configure the SRLG definitions.

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

4. Configure MPLS and the LSPs.

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

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

```
user@PE1 set protocols mpls exclude-srlg
```

6. Enable RSVP on the interfaces.

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

Results

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

```
user@PE1# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
```



```

        family inet {
            address 192.168.12.1/24;
        }
        family mpls;
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 192.168.13.1/24;
        }
        family mpls;
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 192.168.14.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.1/32;
        }
    }
}
}

```

```

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

```

```

user@PE1# show protocols mpls
optimize-timer 120;
label-switched-path pe1-pe2 {

```



```

to 10.255.0.7;
primary via-p1;
secondary path2 {
    standby;
}
}
path via-p1 {
    10.255.0.2 strict;
}
path path2;
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

```

```

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

```

```

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

```

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

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

Verification

Confirm that the configuration is working properly.

Verifying the Secondary Path Link for the LSP

Purpose

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

Action

user@PE1> show mpls lsp detail

```
Ingress LSP: 1 sessions

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

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

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

Meaning

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

RELATED DOCUMENTATION

[SRLG Overview | 277](#)[Example: Configuring SRLG | 278](#)[exclude-srlg | 2608](#)

Example: Configuring SRLG with Link Protection

IN THIS SECTION

- [Requirements | 298](#)
- [Overview | 298](#)
- [Configuration | 299](#)
- [Verification | 323](#)

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

Requirements

This example uses the following hardware and software components:


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

Overview

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

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

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



Configuration

CLI Quick Configuration

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

Router PE1

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.13.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.14.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.1/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
```



```

set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls optimize-timer 120
set protocols mpls label-switched-path pe1-pe2 to 10.255.0.7
set protocols mpls label-switched-path pe1-pe2 link-protection
set protocols mpls label-switched-path pe1-pe2 primary via-p1
set protocols mpls label-switched-path pe1-pe2 secondary path2 standby
set protocols mpls path via-p1 10.255.0.2 strict
set protocols mpls path path2
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P1

```

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

```



```

set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P2

```

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

```

Router P3

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.14.4/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.45.4/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.4/32
set routing-options srlg srlg-a srlg-value 101

```



```

set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P4

```

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

```

Router P5

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.56.6/24
set interfaces ge-0/0/1 unit 0 family mpls

```



```

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

```

Router PE2

```

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

```

Configuring Device PE1

Step-by-Step Procedure

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

To configure the ingress router PE1:

1. Configure the device interfaces.

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

2. Configure OSPF on the interfaces.

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

3. Configure the SRLG definitions.

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

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

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



```
user@PE1# set path via-p1 10.255.0.2 strict
user@PE1# set path path2
```

5. Enable RSVP on the interfaces.

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

Results

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

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



```

        family inet {
            address 10.255.0.1/32;
        }
    }
}
}

```

```

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

```

```

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

```

```

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

```

```

user@PE1# show routing-options
srng {

```



```

srlg-a {
    srlg-value 101;
    srlg-cost 10;
}

```

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

Configuring Device P1

Step-by-Step Procedure

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

To configure device P1:

1. Configure the device interfaces.

[edit interfaces]

```

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

```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```

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

```

3. Configure the SRLG definitions.

[edit routing-options]

```

user@P1# set srlg srlg-a srlg-value 101

```



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

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

```
[edit protocols mpls]
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 srlg srlg-a
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

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

```
[edit protocols rsvp]
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 link-protection
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

Results

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

```
user@P1# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.12.2/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.27.2/24;
    }
    family mpls;
  }
}
```



```

}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 192.168.23.2/24;
    }
    family mpls;
  }
}
ge-0/0/4 {
  unit 0 {
    family inet {
      address 192.168.25.2/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.2/32;
    }
  }
}
}

```

```

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

```

```

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

```



```

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

```

```

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

```

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

Configuring Device P2

Step-by-Step Procedure

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

To configure P2:

1. Configure the device interfaces.

[edit interfaces]

```

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

```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```

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

```



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

3. Configure the SRLG definitions.

```
[edit routing-options]
user@P2# set srlg srlg-a srlg-value 101
user@P2# set srlg srlg-a srlg-cost 10
```

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

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

5. Enable RSVP on the interfaces.

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

Results

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

```
user@P2# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.13.3/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
```



```

    unit 0 {
        family inet {
            address 192.168.37.3/24;
        }
        family mpls;
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 192.168.23.3/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.3/32;
        }
    }
}
}
}

```

```

user@P2# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
    interface ge-0/0/3.0;
    interface lo0.0;
}

```

```

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

```

```

user@P2# show protocols rsvp

```



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

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

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

Configuring Device P3

Step-by-Step Procedure

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

To configure P3:

1. Configure the device interfaces.

[edit interfaces]

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

2. Configure OSPF on the interfaces.

[edit protocols ospf]

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

3. Configure the SRLG definitions.

[edit routing-options]


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

4. Configure MPLS on the interfaces.

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

5. Enable RSVP on the interfaces.

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

Results

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

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



```

        family inet {
            address 10.255.0.4/32;
        }
    }
}

```

```

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

```

```

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

```

```

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

```

```

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

```

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

Configuring Device P4

Step-by-Step Procedure

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

To configure P4:

1. Configure the device interfaces.

[edit interfaces]

```

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

```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```

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

```

3. Configure the SRLG definitions.

[edit routing-options]

```

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

```

4. Configure MPLS on the interfaces.

[edit protocols mpls]

```

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

```

5. Enable RSVP on the interfaces.

[edit protocols rsvp]

```

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

```

Results

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

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

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



```
}
```

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

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

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

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

Configuring Device P5

Step-by-Step Procedure

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

To configure P5:

1. Configure the device interfaces.

[edit interfaces]

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

2. Configure OSPF on the interfaces.


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

3. Configure the SRLG definitions.

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

4. Configure MPLS on the interfaces.

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

5. Enable RSVP on the interfaces.

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

Results

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

```
user@P5# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.56.6/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
```



```

    unit 0 {
        family inet {
            address 192.168.67.6/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.6/32;
        }
    }
}

```

```

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

```

```

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

```

```

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

```

```

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

```

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

Configuring Device PE2

Step-by-Step Procedure

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

To configure PE2:

1. Configure the device interfaces.

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

2. Configure OSPF on the interfaces.

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

3. Configure the SRLG definitions.

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

4. Configure MPLS on the interfaces.

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

5. Enable RSVP on the interfaces.


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

Results

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

```
user@PE2# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.27.7/24;
      }
      family mpls;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 192.168.37.7/24;
      }
      family mpls;
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 192.168.67.7/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.0.7/32;
      }
    }
  }
}
```



```
}
```

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

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

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

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

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

Verification

Confirm that the configuration is working properly.

Verifying the SRLG Cost Is Added to the TE Link

Purpose

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

Action

user@P1> **show ted link detail**

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

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

user@P1> **show rsvp session extensive bypass**

```
Ingress RSVP: 1 sessions

10.255.0.7
  From: 10.255.0.2, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->192.168.27.7
  LSPTtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299776
  Resv style: 1 SE, Label in: -, Label out: 299776
  Time left:    -, Since: Fri Oct 21 13:19:21 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 52081 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
```



```

Path MTU: received 1500
PATH sentto: 192.168.25.5 (ge-0/0/4.0) 26 pkts
RESV rcvfrom: 192.168.25.5 (ge-0/0/4.0) 26 pkts
Explot route: 192.168.25.5 192.168.56.6 192.168.67.7
Record route: <self> 192.168.25.5 192.168.56.6 192.168.67.7
Total 1 displayed, Up 1, Down 0

```

Meaning

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

RELATED DOCUMENTATION

[SRLG Overview | 277](#)

[Example: Configuring SRLG | 278](#)

[Example: Configuring SRLG with Link Protection with the exclude-srlg Option | 325](#)

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

IN THIS SECTION

- [Requirements | 326](#)
- [Overview | 326](#)
- [Configuration | 326](#)
- [Verification | 351](#)

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

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

Router PE1

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

Router P1

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.27.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.23.2/24
```



```

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

```

Router P2

```

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

```



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

Router P3

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

Router P4

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



```

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

```

Router P5

```

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

```

Router PE2

```

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

```



```

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

```

Configuring Device PE1

Step-by-Step Procedure

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

To configure the ingress router PE1:

1. Configure the device interfaces.

```

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

```

2. Configure OSPF on the interfaces.

```

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

```

3. Configure the SRLG definitions.


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

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

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

5. Enable RSVP on the interfaces.

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

Results

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

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



```

unit 0 {
    family inet {
        address 192.168.13.1/24;
    }
    family mpls;
}
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 192.168.14.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.1/32;
        }
    }
}
}
}

```

```

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

```

```

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

```



```

path via-p1 {
    10.255.0.2 strict;
}
path path2;
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

```

```

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

```

```

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

```

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

Configuring Device P1

Step-by-Step Procedure

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

To configure device P1:

1. Configure the device interfaces.

```

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

```


2. Configure OSPF on the interfaces.

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

3. Configure the SRLG definitions.

```
[edit routing-options]
user@P1# set routing-options srlg srlg-a srlg-value 101
user@P1# set routing-options srlg srlg-a srlg-cost 10
```

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

```
[edit protocols mpls]
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 srlg srlg-a
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

5. Enable RSVP on the interfaces and include the **link-protection** statement with the **exclude-srlg** option for interface **ge-0/0/2.0**.

```
[edit protocols rsvp]
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 link-protection exclude-srlg
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

Results

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

```
user@P1# show interfaces
```



```

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

```

```

user@P1# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;

```



```

interface ge-0/0/2.0;
interface ge-0/0/3.0;
interface ge-0/0/4.0;
interface lo0.0;
}

```

```

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

```

```

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

```

```

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

```

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

Configuring Device P2

Step-by-Step Procedure

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

To configure P2:

1. Configure the device interfaces.

[edit interfaces]

```

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

```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```

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

```

3. Configure the SRLG definitions.

[edit routing-options]

```

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

```

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

[edit protocols mpls]

```

user@P2# set interface ge-0/0/1.0
user@P2# set interface ge-0/0/2.0 srlg srlg-a
user@P2# set interface ge-0/0/3.0

```

5. Enable RSVP on the interfaces.

[edit protocols rsvp]

```

user@P2# set interface ge-0/0/1.0
user@P2# set interface ge-0/0/2.0
user@P2# set interface ge-0/0/3.0

```


Results

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

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

```
user@P2# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
```



```

interface ge-0/0/2.0;
interface ge-0/0/3.0;
interface lo0.0;
}

```

```

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

```

```

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

```

```

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

```

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

Configuring Device P3

Step-by-Step Procedure

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

To configure P3:

1. Configure the device interfaces.

```

[edit interfaces]
user@P3# set ge-0/0/1 unit 0 family inet address 192.168.14.4/24
user@P3# set ge-0/0/1 unit 0 family mpls
user@P3# set ge-0/0/2 unit 0 family inet address 192.168.45.4/24

```



```

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

```

2. Configure OSPF on the interfaces.

```

[edit protocols ospf]
user@P3# set traffic-engineering
user@P3# set area 0.0.0.0 interface ge-0/0/1.0
user@P3# set area 0.0.0.0 interface ge-0/0/2.0
user@P3# set area 0.0.0.0 interface lo0.0

```

3. Configure the SRLG definitions.

```

[edit routing-options]
user@P3# set routing-options srlg srlg-a srlg-value 101
user@P3# set routing-options srlg srlg-a srlg-cost 10

```

4. Configure MPLS on the interfaces.

```

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

```

5. Enable RSVP on the interfaces.

```

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

```

Results

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

```

user@P3# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {

```



```

        family inet {
            address 192.168.14.4/24;
        }
        family mpls;
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 192.168.45.4/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.4/32;
        }
    }
}
}
}

```

```

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

```

```

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

```

```

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

```

```

user@P3# show routing-options
srng {

```



```

srlg-a {
    srlg-value 101;
    srlg-cost 10;
}

```

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

Configuring Device P4

Step-by-Step Procedure

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

To configure P4:

1. Configure the device interfaces.

[edit interfaces]

```

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

```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```

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

```

3. Configure the SRLG definitions.

[edit routing-options]

```

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

```


4. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@P4# set interface ge-0/0/1.0
user@P4# set interface ge-0/0/2.0
user@P4# set interface ge-0/0/3.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@P4# set interface ge-0/0/1.0
user@P4# set interface ge-0/0/2.0
user@P4# set interface ge-0/0/3.0
```

Results

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

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



```

    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.5/32;
        }
    }
}
}

```

```

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

```

```

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

```

```

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

```

```

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

```

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

Configuring Device P5

Step-by-Step Procedure

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

To configure P5:

1. Configure the device interfaces.

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

2. Configure OSPF on the interfaces.

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

3. Configure the SRLG definitions.

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

4. Configure MPLS on the interfaces.

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

5. Enable RSVP on the interfaces.

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

Results

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

```
user@P5# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.56.6/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.67.6/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.6/32;
    }
  }
}
```

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

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

```
user@P5# show protocols rsvp
```



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

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

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

Configuring Device PE2

Step-by-Step Procedure

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

To configure PE2:

1. Configure the device interfaces.

[edit interfaces]

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

2. Configure OSPF on the interfaces.

[edit protocols ospf]

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

3. Configure the SRLG definitions.


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

4. Configure MPLS on the interfaces.

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

5. Enable RSVP on the interfaces.

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

Results

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

```
user@PE2# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.27.7/24;
      }
      family mpls;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 192.168.37.7/24;
      }
      family mpls;
    }
  }
}
```



```

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

```

```

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

```

```

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

```

```

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

```

```

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

```



```

    srlg-cost 10;
  }
}

```

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

Verification

Confirm that the configuration is working properly.

Verifying the SRLG Cost Is Added to the TE Link

Purpose

Verify that the TE link is excluded if it belongs to the SRLG of the protected link when **link-protection** is configured with **exclude-srlg**. Issue the **show ted link detail** and **show rsvp session extensive bypass** commands on device P1.

Action

user@P1> **show ted link detail**

```

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

```

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

```

user@P1> **show rsvp session extensive bypass**


```

Ingress RSVP: 1 sessions

10.255.0.7
  From: 10.255.0.2, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->192.168.27.7
  LSPTYPE: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299776
  Resv style: 1 SE, Label in: -, Label out: 299776
  Time left:      -, Since: Fri Oct 21 13:19:21 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 52081 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 192.168.25.5 (ge-0/0/4.0) 63 pkts
  RESV rcvfrom: 192.168.25.5 (ge-0/0/4.0) 63 pkts
  Explct route: 192.168.25.5 192.168.56.6 192.168.67.7
  Record route: <self> 192.168.25.5 192.168.56.6 192.168.67.7
Total 1 displayed, Up 1, Down 0

```

Meaning

The shortest path for the bypass protecting the link P1>PE2 would have been P1>P2>PE2. Because the links P1>PE2 and P2>PE2 both belong to SRLG **srlg-a**, the link P2>PE2 is rejected for CSPF consideration due to the **exclude-srlg** constraint. Therefore, the computed path for the bypass becomes P1>P4>P5>PE2.

RELATED DOCUMENTATION

[SRLG Overview | 277](#)

[Example: Configuring SRLG | 278](#)

[Example: Configuring SRLG with Link Protection | 298](#)

[exclude-srlg | 2608](#)

Configuring MPLS Tunnels

IN THIS CHAPTER

- [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks | 353](#)
- [Configuring IPv6 Tunneling for MPLS | 365](#)
- [Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels | 367](#)
- [Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview | 385](#)
- [Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels | 388](#)
- [Next-Hop-Based Dynamic Tunnel Localization Overview | 401](#)

Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

IN THIS SECTION

- [Requirements | 353](#)
- [Overview | 354](#)
- [Configuration | 356](#)
- [Verification | 364](#)

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

Requirements

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:

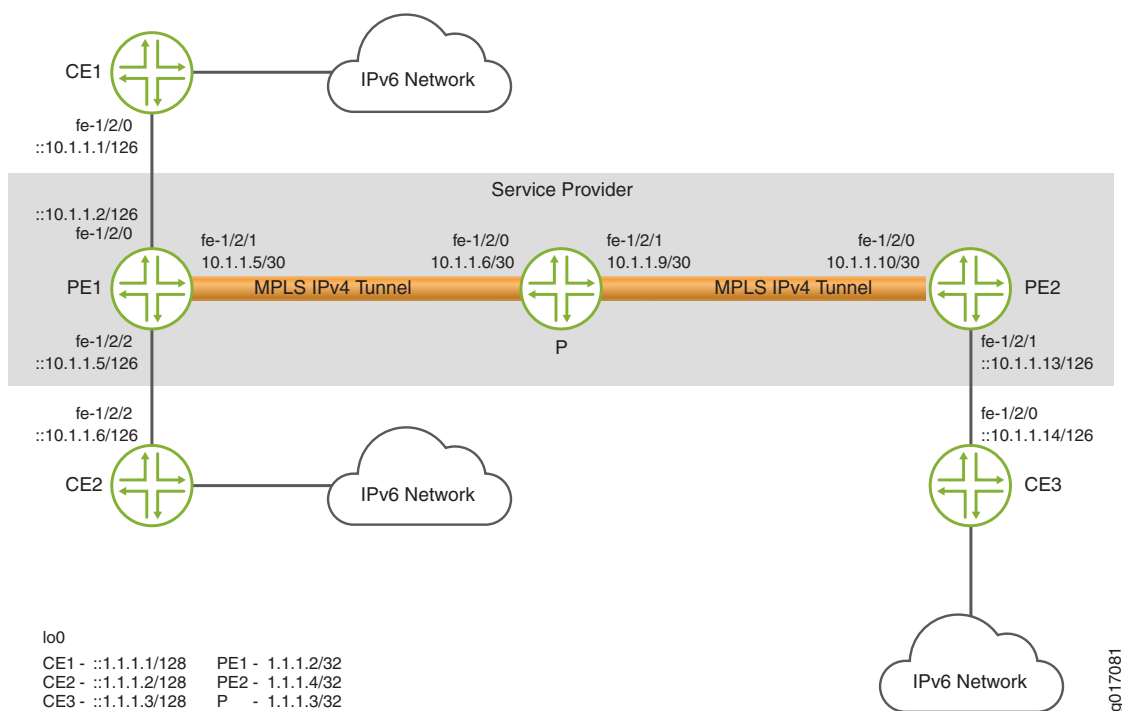
- Internet draft draft-ietf-l3vpn-bgp-ipv6-07.txt, *BGP-MPLS IP VPN extension for IPv6 VPN* (expires January 2006)
- Internet draft draft-ooms-v6ops-bgp-tunnel-06.txt, *Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers* (expires July 2006)

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 17 on page 354](#), Routers PE1 and PE2 are dual-stack BGP routers, meaning they have both IPv4 and IPv6 stacks. The PE routers link the IPv6 networks through the customer edge (CE) routers to the IPv4 core network. The CE routers and the PE routers connect through a link layer that can carry IPv6 traffic. The PE routers use IPv6 on the CE router-facing interfaces and use IPv4 and MPLS on the core-facing interfaces. Note that one of the connected IPv6 networks could be the global IPv6 Internet.

Figure 17: IPv6 Networks Linked by MPLS IPv4 Tunnels



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

The PE routers can learn the IPv6 routes from the CE routers connected to them using routing protocols Routing Information Protocol next generation (RIPng) or MP-BGP, or through static configuration. Note that if BGP is used as the PE-router-to-CE-router protocol, the MP-BGP session between the PE router and CE router could occur over an IPv4 or IPv6 Transmission Control Protocol (TCP) session. Also, the BGP routes exchanged on that session would have SAFI unicast. You must configure an export policy to pass routes between IBGP and EBGP, and between BGP and any other protocol.

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

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

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

When the PE1 router in [Figure 17 on page 354](#) receives an IPv6 packet from the CE1 router, it performs a lookup in the IPv6 forwarding table. If the destination matches a prefix learned from the CE2 router, then no labels need to be pushed and the packet is simply sent to the CE2 router. If the destination matches a prefix that was learned from the PE2 router, then the PE1 router pushes two labels onto the packet and sends it to the provider router. The inner label is 2 and the outer label is the LSP label for the PE2 router.

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

This example includes the following settings:

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

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

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

Configuration

IN THIS SECTION

- [Configuring Device PE1 | 360](#)

CLI Quick Configuration

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

Device PE1


```

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

```

Device PE2

```

set interfaces fe-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces fe-1/2/0 unit 10 family inet6
set interfaces fe-1/2/0 unit 10 family mpls

```



```

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

```

Device P

```

set interfaces fe-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces fe-1/2/0 unit 6 family inet6
set interfaces fe-1/2/0 unit 6 family mpls
set interfaces fe-1/2/1 unit 9 family inet address 10.1.1.9/30
set interfaces fe-1/2/1 unit 9 family inet6
set interfaces fe-1/2/1 unit 9 family mpls

```



```

set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface fe-1/2/0.6
set protocols mpls interface fe-1/2/1.9
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
set protocols ospf area 0.0.0.0 interface fe-1/2/1.9
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ldp interface fe-1/2/0.6
set protocols ldp interface fe-1/2/1.9
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 2

```

Device CE1

```

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

```

Device CE3

```

set interfaces fe-1/2/0 unit 14 family inet6 address ::10.1.1.14/126
set interfaces lo0 unit 5 family inet6 address ::1.1.1.5/128
set protocols bgp group toPE2 type external
set protocols bgp group toPE2 local-address ::10.1.1.14
set protocols bgp group toPE2 family inet6 unicast
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 peer-as 2
set protocols bgp group toPE2 neighbor ::10.1.1.13

```



```

set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 3

```

Configuring Device PE1

Step-by-Step Procedure

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

To configure Device PE1:

1. Configure the interfaces.

```

[edit interfaces]
user@PE1# set fe-1/2/0 unit 2 family inet6 address ::10.1.1.2/126
user@PE1# set fe-1/2/0 unit 2 family mpls
user@PE1# set fe-1/2/1 unit 5 family inet address 10.1.1.5/30
user@PE1# set fe-1/2/1 unit 5 family inet6
user@PE1# set fe-1/2/1 unit 5 family mpls
user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32

```

2. Configure MPLS on the interfaces.

```

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

```

3. Configure BGP.

```

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

```



```

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

```

4. Configure OSPF

```

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

```

5. Configure a signaling protocol.

```

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

```

6. Configure the routing policies.

```

[edit policy-options]
user@PE1# set policy-statement next-hop-self then next-hop self
user@PE1# set policy-statement send-bgp6 from family inet6
user@PE1# set policy-statement send-bgp6 from protocol bgp
user@PE1# set policy-statement send-bgp6 then accept
user@PE1# set policy-statement send-v6 from family inet6
user@PE1# set policy-statement send-v6 from protocol bgp
user@PE1# set policy-statement send-v6 from protocol direct
user@PE1# set policy-statement send-v6 then accept

```

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

```

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

```

Results

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

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

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



```

from {
    family inet6;
    protocol [ bgp direct ];
}
then accept;
}

```

```

user@R1# show protocols
mpls {
    ipv6-tunneling;
    interface fe-1/2/0.2;
    interface fe-1/2/1.5;
}
bgp {
    group toCE1 {
        type external;
        local-address ::10.1.1.2;
        family inet6 {
            unicast;
        }
        export send-bgp6;
        peer-as 1;
        neighbor ::10.1.1.1;
    }
    group toPE2 {
        type internal;
        local-address 1.1.1.2;
        family inet6 {
            labeled-unicast {
                explicit-null;
            }
        }
        export [ next-hop-self send-v6 ];
        neighbor 1.1.1.4;
    }
}
ospf {
    area 0.0.0.0 {
        interface fe-1/2/1.5;
        interface lo0.2 {
            passive;
        }
    }
}
}

```



```
ldp {
  interface fe-1/2/1.5;
}
```

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

If you are done configuring the device, enter **commit** from configuration mode.
Configure the other devices in the topology, as shown in [“CLI Quick Configuration” on page 356](#).

Verification

IN THIS SECTION

- [Verifying That the CE Devices Have Connectivity | 364](#)

Confirm that the configuration is working properly.

Verifying That the CE Devices Have Connectivity

Purpose

Make sure that the tunnel is operating.

Action

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

```
user@CE1> ping ::10.1.1.14
```

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

```
user@CE3> ping ::10.1.1.1
```

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



```
16 bytes from ::10.1.1.1, icmp_seq=1 hlim=61 time=1.338 ms
16 bytes from ::10.1.1.1, icmp_seq=2 hlim=61 time=1.351 ms
```

Meaning

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

RELATED DOCUMENTATION

[Configuring the Ingress Router for MPLS-Signaled LSPs | 516](#)

[Minimum RSVP Configuration | 827](#)

Configuring IPv6 Tunneling for MPLS

You can configure the IPv6 tunneling for MPLS to tunnel IPv6 traffic over an MPLS-based IPv4 network. This configuration allows you to interconnect a number of smaller IPv6 networks over an IPv4-based network core, giving you the ability to provide IPv6 service without having to upgrade the switches in your core network. BGP is configured to exchange routes between the IPv6 networks, and data is tunneled between these IPv6 networks by means of IPv4-based MPLS.

To configure IPv6 tunneling for MPLS on your EX Series switch:

1. Configure IPv4 and IPv6 IP addresses for all the core interfaces:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
```

2. Configure the number assigned to you by the Network Information Center (NIC) as the autonomous system (AS) number

```
[edit routing-options]
user@switch# set autonomous-system number
```

3. Advertise label 0 to the egress router of the LSP:

```
[edit protocols]
user@switch# set mpls explicit-null
```

4. Configure the LSP to allow IPv6 routes to be resolved over an MPLS network by converting all routes stored in the inet3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table:


```
[edit protocols]
user@switch# set mpls ipv6-tunneling
```

5. Set the local AS number:

```
[edit protocols bgp]
user@switch# set local-as local-autonomous-system-number
```

6. Configure the default import and export policies:

```
[edit protocols bgp]
user@switch# set local-address address
user@switch# set import default-import
user@switch# set family inet6 labeled-unicast explicit-null
user@switch# set export default-export
```

7. Configure a BGP group that recognizes only the specified BGP systems as peers. Define a group name, group type, local end of a BGP session, and a neighbor (peer). To configure multiple BGP peers, include multiple neighbor statements:

```
[edit protocols bgp]
user@switch# set group group-name type internal
user@switch# set group group-name local-address address-of-the-local-end-of-a-bgp-session
user@switch# set group group-name family inet6 labeled-unicast explicit-null
user@switch# set group group-name peer-as peer-autonomous-system-number
user@switch# set group group-name neighbor address family inet6 labeled-unicast explicit-null
```

8. Configure routing options to accept the default import and export policies:

```
[edit policy-options]
user@switch# set policy-statement default-import then accept
user@switch# set policy-statement default-export then accept
```

RELATED DOCUMENTATION

| [Example: Configuring MPLS on EX8200 and EX4500 Switches](#) | 59

Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels

IN THIS SECTION

- [Requirements | 367](#)
- [Overview | 368](#)
- [Configuration | 371](#)
- [Verification | 379](#)
- [Troubleshooting | 384](#)

This example shows how to configure a dynamic MPLS-over-UDP tunnel that includes a tunnel composite next hop. The MPLS-over-UDP feature provides a scaling advantage on the number of IP tunnels supported on a device.

Starting in Junos OS Release 18.3R1, MPLS-over-UDP tunnels are supported on PTX Series routers and QFX Series switches. For every dynamic tunnel configured on a PTX router or a QFX switch, a tunnel composite next hop, an indirect next hop, and a forwarding next hop is created to resolve the tunnel destination route. You can also use policy control to resolve the dynamic tunnel over select prefixes by including the [forwarding-rib](#) configuration statement at the **[edit routing-options dynamic-tunnels]** hierarchy level.

Requirements

This example uses the following hardware and software components:

- Five MX Series routers with MPCs and MICs.
- Junos OS Release 16.2 or later running on the PE routers.

Before you begin:

1. Configure the device interfaces, including the loopback interface.
2. Configure the router ID and autonomous system number for the device.
3. Establish an internal BGP (IBGP) session with the remote PE device.
4. Establish OSPF peering among the devices.

Overview

Starting with Junos OS Release 16.2, a dynamic UDP tunnel supports the creation of a tunnel composite next hop for every UDP tunnel configured. These next-hop-based dynamic UDP tunnels are referred to as MPLS-over-UDP tunnels. The tunnel composite next hop are enabled by default for the MPLS-over-UDP tunnels.

MPLS-over-UDP tunnels can be bidirectional or unidirectional in nature. When the PE devices are connected over MPLS-over-UDP tunnels in both directions, it is called a bidirectional MPLS-over-UDP tunnel. When two PE devices are connected over MPLS-over-UDP tunnel in one direction, and over MPLS/IGP in the other direction, it is called an unidirectional MPLS-over-UDP tunnel.

Unidirectional MPLS-over-UDP tunnels are used in migration scenarios, or in cases where two PE devices provide connectivity to each other over two disjoint networks. Because reverse direction tunnel does not exist for unidirectional MPLS-over-UDP tunnels, you must configure a filter-based MPLS-over-UDP decapsulation on the remote PE device for forwarding the traffic.

Starting in Junos OS Release 18.2R1, on PTX series routers and QFX10000 with unidirectional MPLS-over-UDP tunnels, you must configure the remote PE device with an input filter for MPLS-over-UDP packets, and an action for decapsulating the IP and UDP headers for forwarding the packets in the reverse tunnel direction.

For example, on the remote PE device, Device PE2, the following configuration is required for unidirectional MPLS-over-UDP tunnels:

PE2

```
[edit firewall filter]
user@host# set Decap_Filter term udp_decap from protocol udp
user@host# set Decap_Filter term udp_decap from destination-port 6635
user@host# set Decap_Filter term udp_decap then count UDP_PKTS
user@host# set Decap_Filter term udp_decap then decapsulate mpls-in-udp
user@host# set Decap_Filter term def then count def_pkt
user@host# set Decap_Filter term def then accept
```

In the above sample configuration, *Decap_Filter* is the name of the firewall filter used for MPLS-over-UDP decapsulation. The term *udp_decap* is the input filter for accepting UDP packets on the core-facing interface of Device PE2, and then decapsulate the MPLS-over-UDP packets to MPLS-over-IP packets for forwarding.

You can use the existing firewall operational mode commands, such as **show firewall filter** to view the filter-based MPLS-over-UDP decapsulation.

For example:

```
user@host >show firewall filter Decap_Filter
```

```
Filter: Decap_Filter
Counters:
Name                               Bytes           Packets
UDP_PKTS                          16744           149
def_pkt                            13049           136
```

NOTE:

For unidirectional MPLS-over-UDP tunnels:

- Only IPv4 address is supported as the outer header. Filter-based MPLS-over-UDP decapsulation does not support IPv6 address in the outer header.
- Only the default routing instance is supported after decapsulation.

Starting in Junos OS Release 17.1, on MX Series routers with MPCs and MICs, the scaling limit of MPLS-over-UDP tunnels is increased.

Starting in Junos Release 19.2R1, on MX Series routers with MPCs and MICs, carrier supporting carrier (CSC) architecture can be deployed with MPLS-over-UDP tunnels carrying MPLS traffic over dynamic IPv4 UDP tunnels that are established between supporting carrier's PE devices. With this enhancement, the scaling advantage that the MPLS-over-UDP tunnels provided is further increased. The CSC support with MPLS-over-UDP tunnel is not supported for IPv6 UDP tunnel.

The existing dynamic tunnel feature requires complete static configuration. Currently, the tunnel information received from peer devices in advertised routes is ignored. Starting in Junos OS Release 17.4R1, on MX Series routers, the next-hop-based dynamic MPLS-over-UDP tunnels are signaled using BGP encapsulation extended community. BGP export policy is used to specify the tunnel types, advertise the sender side tunnel information, and parse and convey the receiver side tunnel information. A tunnel is created according to the received type tunnel community.

Multiple tunnel encapsulations are supported by BGP. On receiving multiple capability, the next-hop-based dynamic tunnel is created based on the configured BGP policy and tunnel preference. The tunnel preference should be consistent across both the tunnel ends for the tunnel to be set up. By default, MPLS-over-UDP tunnel is preferred over GRE tunnels. If dynamic tunnel configuration exists, it takes precedence over received tunnel community.

When configuring a next-hop-based dynamic MPLS-over-UDP tunnel, be aware of the following considerations:

- An IBGP session must be configured between the PE devices.
- A switchover between the next-hop-based dynamic tunnel encapsulations (UDP and GRE) is allowed, and this can impact network performance in terms of the supported IP tunnel scaling values in each mode.
- Having both GRE and UDP next-hop-based dynamic tunnel encapsulation types for the same tunnel destination leads to a commit failure.
- For unidirectional MPLS-over-UDP tunnels, you must explicitly configure filter-based MPLS-over-UDP decapsulation on the remote PE device for the packets to be forwarded.
- Graceful Routing Engine switchover (GRES) is supported with MPLS-over-UDP, and the MPLS-over-UDP tunnel type flags are unified ISSU and NSR compliant.
- MPLS-over-UDP tunnels are supported on virtual MX (vMX).
- MPLS-over-UDP tunnels support dynamic GRE tunnel creation based upon new IPv4-mapped-IPv6 next hops.
- MPLS-over-UDP tunnel are supported in interoperability with contrail, wherein the MPLS-over-UDP tunnels are created from the contrail vRouter to an MX gateway. To enable this, the following community is required to be advertised in the route from the MX Series router to the contrail vRouter:

```
[edit policy-options community]
udp members 0x030c:64512:13;
```


At a given point in time, only one tunnel type is supported on the contrail vRouter—next-hop-based dynamic GRE tunnels, MPLS-over-UDP tunnels, or VXLAN.

- The following features are not supported with the next-hop-based dynamic MPLS-over-UDP tunnel configuration:
 - RSVP automatic mesh
 - Plain IPV6 GRE and UDP tunnel configuration
 - Logical systems

Topology

[Figure 18 on page 371](#) illustrates a Layer 3 VPN scenario over dynamic MPLS-over-UDP tunnels. The customer edge (CE) devices, CE1 and CE2, connect to provider edge (PE) devices, PE1 and PE2, respectively. The PE devices are connected to a provider device (Device P1), and an internal BGP (IBGP) session interconnects the two PE devices. A dynamic next-hop-based bidirectional MPL-over-UDP tunnel is configured between the PE devices.

Figure 18: Dynamic MPLS-over-UDP Tunnels



The MPLS-over-UDP tunnel is handled as follows:

1. After a MPLS-over-UDP tunnel is configured, a tunnel destination mask route with a tunnel composite next hop is created for the tunnel in the `inet.3` routing table. This IP tunnel route is withdrawn only when the dynamic tunnel configuration is deleted.

The tunnel composite next-hop attributes include the following:

- When Layer 3 VPN composite next hop is disabled—Source and destination address, encapsulation string, and VPN label.
 - When Layer 3 VPN composite next hop and per-prefix VPN label allocation are enabled—Source address, destination address, and encapsulation string.
 - When Layer 3 VPN composite next hop is enabled and per-prefix VPN label allocation is disabled—Source address, destination address, and encapsulation string. The route in this case is added to the other virtual routing and forwarding instance table with a secondary route.
2. The PE devices are interconnected using an IBGP session. The IBGP route next hop to a remote BGP neighbor is the protocol next hop, which is resolved using the tunnel mask route with the tunnel next hop.
 3. After the protocol next hop is resolved over the tunnel composite next hop, indirect next hops with forwarding next hops are created.
 4. The tunnel composite next hop is used to forward the next hops of the indirect next hops.

Configuration

CLI Quick Configuration

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

CE1

```
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.1/8
set interfaces lo0 unit 0 family inet address 127.0.0.1/8
set routing-options router-id 127.0.0.1
set routing-options autonomous-system 200
set protocols bgp group ce1-pe1 export export-loopback-direct
set protocols bgp group ce1-pe1 peer-as 100
set protocols bgp group ce1-pe1 neighbor 10.0.0.2
set policy-options policy-statement export-loopback-direct term term-1 from interface lo0.0
set policy-options policy-statement export-loopback-direct term term-1 from route-filter 127.0.0.1/8
  exact
set policy-options policy-statement export-loopback-direct term term-1 then accept
```

CE2

```
set interfaces ge-0/0/0 unit 0 family inet address 203.0.113.2/24
set interfaces lo0 unit 0 family inet address 127.0.0.5/8
set routing-options router-id 127.0.0.5
set routing-options autonomous-system 200
set protocols bgp group ce1-pe1 export export-loopback-direct
set protocols bgp group ce1-pe1 peer-as 100
set protocols bgp group ce1-pe1 neighbor 203.0.113.1
set policy-options policy-statement export-loopback-direct term term-1 from interface lo0.0
set policy-options policy-statement export-loopback-direct term term-1 from route-filter 127.0.0.5/8
  exact
set policy-options policy-statement export-loopback-direct term term-1 then accept
```

PE1

```
set interfaces ge-0/0/0 unit 0 family inet address 10.0.0.2/8
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/1 unit 0 family mpls
```



```

set interfaces lo0 unit 0 family inet address 127.0.0.2/8
set routing-options static route 33.0.0.0/8 next-hop 192.0.2.2
set routing-options router-id 127.0.0.2
set routing-options autonomous-system 100
set routing-options forwarding-table export pplb
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe2 source-address 127.0.0.2
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe2 udp
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe2 destination-networks 127.0.0.0/8
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 127.0.0.2
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 127.0.0.4
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-instances MPLS-over-UDP-PE1 instance-type vrf
set routing-instances MPLS-over-UDP-PE1 interface ge-0/0/0.0
set routing-instances MPLS-over-UDP-PE1 route-distinguisher 127.0.0.2:1
set routing-instances MPLS-over-UDP-PE1 vrf-target target:600:1
set routing-instances MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 peer-as 200
set routing-instances MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 neighbor 10.0.0.1 as-override

```

P1

```

set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.3/8
set routing-options router-id 127.0.0.3
set routing-options autonomous-system 100
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

PE2


```

set interfaces ge-0/0/0 unit 0 family inet address 203.0.113.1/24
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.4/8
set routing-options nonstop-routing
set routing-options router-id 127.0.0.4
set routing-options autonomous-system 100
set routing-options forwarding-table export pplib
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe1 source-address 127.0.0.4
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe1 udp
set routing-options dynamic-tunnels udp-dyn-tunnel-to-pe1 destination-networks 127.0.0.0/8
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 127.0.0.4
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 127.0.0.2
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set routing-instances MPLS-over-UDP-PE2 instance-type vrf
set routing-instances MPLS-over-UDP-PE2 interface ge-0/0/0.0
set routing-instances MPLS-over-UDP-PE2 route-distinguisher 127.0.0.4:1
set routing-instances MPLS-over-UDP-PE2 vrf-target target:600:1
set routing-instances MPLS-over-UDP-PE2 protocols bgp group ebgp peer-as 200
set routing-instances MPLS-over-UDP-PE2 protocols bgp group ebgp neighbor 203.0.113.2 as-override

```

Step-by-Step Procedure

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

To configure Device PE1:

1. Configure the device interfaces including the loopback interface of the device.

```

[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 10.0.0.2/8
user@PE1# set ge-0/0/1 unit 0 family inet address 192.0.2.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 127.0.0.2/8

```

2. Configure a static route for routes from Device PE1 with Device P1 as the next-hop destination.


```
[edit routing-options]
user@PE1# set static route 33.0.0.0/8 next-hop 192.0.2.2
```

3. Configure the router-ID and autonomous system number for Device PE1.

```
[edit routing-options]
user@PE1# set router-id 127.0.0.2
user@PE1# set autonomous-system 100
```

4. (PTX Series only) Configure policy control to resolve the MPLS-over-UDP dynamic tunnel route over select prefixes.

```
[edit routing-options dynamic-tunnels]
user@PTX-PE1# set forwarding-rib inet.0 inet-import dynamic-tunnel-fwd-route-import
```

5. (PTX Series only) Configure the inet-import policy for resolving dynamic tunnel destination routes over

```
[edit policy-options]
user@PTX-PE1# set policy-statement dynamic-tunnel-fwd-route-import term 1 from route-filter 127.0.0.0/8
exact
user@PTX-PE1# set policy-statement dynamic-tunnel-fwd-route-import term 1 then accept
user@PTX-PE1# set policy-options policy-statement dynamic-tunnel-fwd-route-import then reject
```

6. Configure IBGP peering between the PE devices.

```
[edit protocols]
user@PE1# set bgp group IBGP type internal
user@PE1# set bgp group IBGP local-address 127.0.0.2
user@PE1# set bgp group IBGP family inet-vpn unicast
user@PE1# set bgp group IBGP neighbor 127.0.0.4
```

7. Configure OSPF on all the interfaces of Device PE1, excluding the management interface.

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


8. Enable next-hop-based dynamic GRE tunnel configuration on Device PE1.

NOTE: This step is required only for illustrating the implementation difference between next-hop-based dynamic GRE tunnels and MPLS-over-UDP tunnels.

```
[edit routing-options]
user@PE1# set dynamic-tunnels gre next-hop-based-tunnel
```

9. Configure the MPLS-over-UDP tunnel parameters from Device PE1 to Device PE2.

```
[edit routing-options]
user@PE1# set dynamic-tunnels udp-dyn-tunnel-to-pe2 source-address 127.0.0.2
user@PE1# set dynamic-tunnels udp-dyn-tunnel-to-pe2 udp
user@PE1# set dynamic-tunnels udp-dyn-tunnel-to-pe2 destination-networks 127.0.0.0/8
```

10. Configure a VRF routing instance on Device PE1 and other routing instance parameters.

```
[edit routing-instances]
user@PE1# set MPLS-over-UDP-PE1 instance-type vrf
user@PE1# set MPLS-over-UDP-PE1 interface ge-0/0/0.0
user@PE1# set MPLS-over-UDP-PE1 route-distinguisher 127.0.0.2:1
user@PE1# set MPLS-over-UDP-PE1 vrf-target target:600:1
```

11. Enable BGP in the routing instance configuration for peering with Device CE1.

```
[edit routing-instances]
user@PE1# set MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 peer-as 200
user@PE1# set MPLS-over-UDP-PE1 protocols bgp group pe1-ce1 neighbor 10.0.0.1 as-override
```

Results

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

```
user@PE1# show interfaces
ge-0/0/0 {
```



```

    unit 0 {
        family inet {
            address 10.0.0.2/8;
        }
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 192.0.2.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.2/8;
        }
    }
}

```

```

user@PE1# show routing-options
static {
    route 33.0.0.0/8 next-hop 192.0.2.2;
}
router-id 127.0.0.2;
autonomous-system 100;
forwarding-table {
    export pplb;
}
dynamic-tunnels {
    gre next-hop-based-tunnel;
    udp-dyn-tunnel-to-pe2 {
        source-address 127.0.0.2;
        udp;
        destination-networks {
            127.0.0.0/8;
        }
    }
}

```

```

user@PE1# show protocols

```



```

bgp {
  group IBGP {
    type internal;
    local-address 127.0.0.2;
    family inet-vpn {
      unicast;
    }
    neighbor 127.0.0.4;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface lo0.0 {
      passive;
    }
  }
}

```

user@PE1# **show routing-instances**

```

MPLS-over-UDP-PE1 {
  instance-type vrf;
  interface ge-0/0/0.0;
  route-distinguisher 127.0.0.2:1;
  vrf-target target:600:1;
  protocols {
    bgp {
      group pe1-ce1 {
        peer-as 200;
        neighbor 10.0.0.1 {
          as-override;
        }
      }
    }
  }
}

```

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

Verification

IN THIS SECTION

- [Verifying the Connection Between PE Devices | 379](#)
- [Verify the Dynamic Tunnel Routes on Device PE1 | 380](#)
- [Verify the Dynamic Tunnel Routes on Device PE2 | 382](#)
- [Verifying That the Routes Have the Expected Indirect-Next-Hop Flag | 382](#)

Confirm that the configuration is working properly.

Verifying the Connection Between PE Devices

Purpose

Verify the BGP peering status between Device PE1 and Device PE2, and the BGP routes received from Device PE2.

Action

From operational mode, run the **show bgp summary** and **show route receive-protocol bgp ip-address table bgp.l3vpn.0** commands.

```
user@PE1> show bgp summary
```

```
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History  Damp State   Pending
bgp.l3vpn.0
                2          2          0          0          0          0
Peer           AS        InPkt    OutPkt    OutQ     Flaps  Last Up/Dwn
State|#Active/Received/Accepted/Damped...
127.0.0.4      100        139      136        0         0      58:23 Establ

  bgp.l3vpn.0: 2/2/2/0
  MPLS-over-UDP-PE1.inet.0: 2/2/2/0
10.0.0.1       200        135      136        0         0      58:53 Establ

  MPLS-over-UDP-PE1.inet.0: 1/1/1/0
```

```
user@PE1> show route receive-protocol bgp 127.0.0.4 table bgp.l3vpn.0
```



```

bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
  127.0.0.4:1:127.0.0.5/8
*                   127.0.0.4                100        200 I
  127.0.0.4:1:200.1.1.0/24
*                   127.0.0.4                100        I

```

Meaning

- In the first output, the BGP session state is **Establ**, which means that the session is up and the PE devices are peered.
- In the second output, Device PE1 has learned two BGP routes from Device PE2.

Verify the Dynamic Tunnel Routes on Device PE1

Purpose

Verify the routes in the inet.3 routing table and the dynamic tunnel database information on Device PE1.

Action

From operational mode, run the **show route table inet.3**, **show dynamic-tunnels database terse**, **show dynamic-tunnels database**, and **show dynamic-tunnels database summary** commands.

```
user@PE1> show route table inet.3
```

```

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.0/8      *[Tunnel/300] 00:21:18
                  Tunnel
127.0.0.4/8      *[Tunnel/300] 00:21:18
                  Tunnel Composite

```

```
user@PE1> show dynamic-tunnels database terse
```

```

Table: inet.3

Destination-network: 127.0.0.0/8

```


Destination	Source	Next-hop	Type	Status
127.0.0.4/8	127.0.0.2	0xb395b10 nhid 613	udp Up	

user@PE1> **show dynamic-tunnels database**

```
Table: inet.3

Destination-network: 55.0.0.0/8

Destination-network: 55.66.0.0/16

Destination-network: 55.66.77.0/24
Tunnel to: 127.0.0.4/8
Reference count: 2
Next-hop type: UDP
Source address: 127.0.0.2 Tunnel Id: 2
Next hop: tunnel-composite, 0xb395b10, nhid 613
VPN Label: Push 299776 Reference count: 3
Traffic Statistics: Packets 0, Bytes 0
State: Up
```

user@PE1> **show dynamic-tunnels database summary**

```
Dynamic Tunnels, Total 1 displayed
GRE Tunnel:
Active Tunnel Mode, Next Hop Base
IFL Based, Total 0 displayed, Up 0, Down 0
Nexthop Based, Total 0 displayed, Up 0, Down 0
RSVP Tunnel:
Total 0 displayed
UDP Tunnel:
Total 1 displayed, Up 1, Down 0
```

Meaning

- In the first output, because Device PE1 is configured with the MPLS-over-UDP tunnel, a tunnel composite route is created for the inet.3 routing table route entry.
- In the remaining outputs, the MPLS-over-UDP tunnel is displayed with the tunnel encapsulation type, tunnel next hop parameters, and tunnel status.

Verify the Dynamic Tunnel Routes on Device PE2

Purpose

Verify the routes in the inet.3 routing table and the dynamic tunnel database information on Device PE2.

Action

From operational mode, run the **show route table inet.3**, and the **show dynamic-tunnels database terse** commands.

user@PE2> **show route table inet.3**

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.0/8      *[Tunnel/300] 00:39:31
                  Tunnel
127.0.0.2/8      *[Tunnel/300] 00:24:53
                  Tunnel Composite
```

user@PE1> **show dynamic-tunnels database terse**

```
Table: inet.3

Destination-network: 127.0.0.0/8
Destination      Source      Next-hop      Type      Status
127.0.0.2/8      127.0.0.4   0xb395450 nhid 615   udp      Up
```

Meaning

The outputs show the MPLS-over-UDP tunnel creation and the next-hop ID assigned as the next-hop interface, similar to Device PE1.

Verifying That the Routes Have the Expected Indirect-Next-Hop Flag

Purpose

Verify that Device PE1 and Device PE2 are configured to maintain the indirect next hop to forwarding next-hop binding on the Packet Forwarding Engine forwarding table.

Action

From operational mode, run the **show krt indirect-next-hop** command on Device PE1 and Device PE2.

user@PE1> **show krt indirect-next-hop**


```

Indirect Nexthop:
Index: 1048574 Protocol next-hop address: 127.0.0.4
  RIB Table: bgp.l3vpn.0
  Label: Push 299776
  Policy Version: 1                      References: 1
  Locks: 3                               0xb2ab630
  Flags: 0x0
  INH Session ID: 0x0
  INH Version ID: 0
  Ref RIB Table: unknown
    Tunnel type: UDP, Reference count: 3, nhid: 613
    Destination address: 127.0.0.4, Source address: 127.0.0.2
    Tunnel id: 2, VPN Label: Push 299776, TTL action: prop-ttl
    IGP FRR Interesting proto count : 1
    Chain IGP FRR Node Num          : 1
      IGP Resolver node(hex)        : 0xb3c70dc
      IGP Route handle(hex)         : 0xb1ae688      IGP rt_entry protocol
: Tunnel
      IGP Actual Route handle(hex) : 0x0            IGP Actual rt_entry protocol
: Any

```

user@PE2> show krt indirect-next-hop

```

Indirect Nexthop:
Index: 1048575 Protocol next-hop address: 127.0.0.2
  RIB Table: bgp.l3vpn.0
  Label: Push 299776
  Policy Version: 1                      References: 2
  Locks: 3                               0xb2ab740
  Flags: 0x0
  INH Session ID: 0x0
  INH Version ID: 0
  Ref RIB Table: unknown
    Tunnel type: UDP, Reference count: 3, nhid: 615
    Destination address: 127.0.0.2, Source address: 127.0.0.4
    Tunnel id: 1, VPN Label: Push 299776, TTL action: prop-ttl
    IGP FRR Interesting proto count : 2
    Chain IGP FRR Node Num          : 1
      IGP Resolver node(hex)        : 0xb3d3a28
      IGP Route handle(hex)         : 0xb1ae634      IGP rt_entry protocol
: Tunnel
      IGP Actual Route handle(hex) : 0x0            IGP Actual rt_entry protocol
: Any

```


Meaning

The outputs show that a next-hop-based dynamic MPLS-over-UDP tunnel is created between the PE devices.

Troubleshooting

IN THIS SECTION

- [Troubleshooting Commands | 384](#)

To troubleshoot the next-hop-based dynamic tunnels, see:

Troubleshooting Commands

Problem

The next-hop-based dynamic MPLS-over-UDP tunnel configuration is not taking effect.

Solution

To troubleshoot the next-hop-based MPLS-over-UDP tunnel configuration, use the following **traceroute** commands at the **[edit routing-options dynamic-tunnels]** statement hierarchy:

- **traceoptions file *file-name***
- **traceoptions file size *file-size***
- **traceoptions flag all**

For example:

```
[edit routing-options dynamic-tunnels]
traceoptions {
  file udp_dyn_pe1.wri size 4294967295;
  flag all;
}
```


Release History Table

Release	Description
19.2R1	Starting in Junos Release 19.2R1, on MX Series routers with MPCs and MICs, carrier supporting carrier (CSC) architecture can be deployed with MPLS-over-UDP tunnels carrying MPLS traffic over dynamic IPv4 UDP tunnels that are established between supporting carrier's PE devices.
18.3R1	Starting in Junos OS Release 18.3R1, MPLS-over-UDP tunnels are supported on PTX Series routers and QFX Series switches.
18.2R1	Starting in Junos OS Release 18.2R1, on PTX series routers and QFX10000 with unidirectional MPLS-over-UDP tunnels, you must configure the remote PE device with an input filter for MPLS-over-UDP packets, and an action for decapsulating the IP and UDP headers for forwarding the packets in the reverse tunnel direction.
17.4R1	Starting in Junos OS Release 17.4R1, on MX Series routers, the next-hop-based dynamic MPLS-over-UDP tunnels are signaled using BGP encapsulation extended community.
17.1R1	Starting in Junos OS Release 17.1, on MX Series routers with MPCs and MICs, the scaling limit of MPLS-over-UDP tunnels is increased.

RELATED DOCUMENTATION

Configuring GRE Tunnels for Layer 3 VPNs

Example: Configuring a Next-Hop-Based Dynamic GRE Tunnels

Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview

With the rise in deployment of high-scale IP tunnels in data centers, there is a need to add security measures that allow users to limit malicious traffic from compromised virtual machines (VMs). One possible attack is the injecting of traffic into an arbitrary customer VPN from a compromised server through the gateway router. In such cases, anti-spoofing checks on IP tunnels ensure that only legitimate sources are injecting traffic into data centers from their designated IP tunnels.


Next-hop-based dynamic IP tunnels create a tunnel composite next hop for every dynamic tunnel created on the device. Because next-hop-based dynamic tunnels remove the dependency on physical interfaces for every new dynamic tunnel configured, configuring next-hop-based dynamic tunnels provides a scaling advantage over the number of dynamic tunnels that can be created on a device. Starting in Junos OS Release 17.1, anti-spoofing capabilities for next-hop-based dynamic IP tunnels is provided for

next-hop-based dynamic tunnels. With this enhancement, a security measure is implemented to prevent injecting of traffic into an arbitrary customer VPN from a compromised server through the gateway router.

Anti-spoofing is implemented using reverse path forwarding checks in the Packet Forwarding Engine. The checks are implemented for the traffic coming through the tunnel to the routing instance. Currently, when the gateway router receives traffic from a tunnel, only the destination lookup is done and the packet is forwarded accordingly. When anti-spoofing protection is enabled, the gateway router also does a source address lookup of the encapsulation packet IP header in the VPN, in addition to the tunnel destination lookup. This ensures that legitimate sources are injecting traffic through their designated IP tunnels. As a result, anti-spoofing protection ensures that the tunnel traffic is received from a legitimate source on the designated tunnels.

Figure 19 on page 386 illustrates a sample topology with the requirements for anti-spoofing protection.

Figure 19: Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels



In this example, the gateway router is Router G. Router G has two VPNs—Green and Blue. The two servers, Server A and Server B, can reach the Green and Blue VPNs on Router G through the next-hop-based dynamic tunnels T1 and T2, respectively. Several hosts and virtual machines (P, Q, R, S, and T) connected to the servers can reach the VPNs through the gateway router, Router G. Router G has the virtual routing and forwarding (VRF) tables for Green and Blue VPNs, each populated with the reachability information for the virtual machines in those VPNs.

For example, in VPN Green, Router G uses tunnel T1 to reach host P, tunnel T2 to reach hosts R and S, and load balancing is done between tunnels T1 and T2 to reach the multihomed host Q. In VPN Blue, Router G uses tunnel T1 to reach hosts P and R, and tunnel T2 to reach hosts Q and T.

The check passes for reverse path forwarding when:

- A packet comes from a legitimate source on its designated tunnel.

Host P in VPN Green sends a packet to host X using tunnel T1. Because Router G can reach host P through tunnel T1, it allows the packet to pass and forwards the packet to host X.

- A packet comes from a multihomed source on its designated tunnels.

Host Q in VPN Green is multihomed on servers A and B, and can reach Router G through tunnels T1 and T2. Host Q sends a packet to host Y using tunnel T1, and a packet to host X using tunnel T2. Because Router G can reach host Q through tunnels T1 and T2, it allows the packets to pass and forwards them to hosts Y and X, respectively.

Layer 3 VPNs do not have anti-spoofing protection enabled by default. To enable anti-spoofing for next-hop-based dynamic tunnels, include the **ip-tunnel-rpf-check** statement at the **[edit routing-instances routing-instance-name routing-options forwarding-table]** hierarchy level. The reverse path forwarding check is applied to the VRF routing instance only. The default mode is set to **strict**, where the packet that comes from a source on a nondesignated tunnel does not pass the check. The **ip-tunnel-rpf-check** mode can be set as **loose**, where the reverse path forwarding check fails when the packet comes from a nonexistent source. An optional firewall filter can be configured under the **ip-tunnel-rpf-check** statement to count and log the packets that failed the reverse path forwarding check.

The following sample output shows an anti-spoofing configuration:

```
[edit routing-instances routing-instance-name routing-options forwarding-table]
ip-tunnel-rpf-check {
  mode loose;
  fail-filter filter-name;
}
```

Take the following guidelines under consideration when configuring anti-spoofing protection for next-hop-based dynamic tunnels:

- Anti-spoofing protection can be enabled for IPv4 tunnels and IPv4 data traffic only. The anti-spoofing capabilities are not supported on IPv6 tunnels and IPv6 data traffic.
- Anti-spoofing for next-hop-based dynamic tunnels can detect and prevent a compromised virtual machine (inner source reverse path forwarding check) but not a compromised server that is label-spoofing.
- The next-hop-based IP tunnels can originate and terminate on an inet.0 routing table.

- Anti-spoofing protection is effective when the VRF routing instance has label-switched interfaces (LSIs) (using the **vrf-table-label**), or virtual tunnel (VT) interfaces. With **per-next-hop** label on the VRF routing instance, anti-spoofing protection is not supported.
- The **rpf fail-filter** is applicable only to the inner IP packet.
- Enabling anti-spoofing checks does not affect the scaling limit of the next-hop-based dynamic tunnels on a device.
- The system resource utilization with anti-spoofing protection enabled for the VRF routing instance is slightly higher than the utilization of next-hop-based dynamic tunnels without the anti-spoofing protection enabled.
- Anti-spoofing protection requires additional source IP address checks, which has minimal impact on network performance.
- Graceful Routing Engine switchover (GRES) and in-service software upgrade (ISSU) are supported with anti-spoofing protection.

RELATED DOCUMENTATION

[ip-tunnel-rpf-check | 2632](#)

[Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels | 388](#)

Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels

IN THIS SECTION

- [Requirements | 389](#)
- [Overview | 389](#)
- [Configuration | 390](#)
- [Verification | 398](#)

This example shows how to configure reverse path forwarding checks for the virtual routing and forwarding (VRF) routing instance to enable anti-spoofing protection for next-hop-based dynamic tunnels. The checks ensure that legitimate sources are injecting traffic through their designated IP tunnels.

Requirements

This example uses the following hardware and software components:

- Three MX Series Routers with MICs, each connected to a host device.
- Junos OS Release 17.1 or later running on one or all the routers.

Before you begin:

- Enable tunnel services configuration on the Flexible PIC Concentrator.
- Configure the router interfaces.
- Configure the router-ID and assign an autonomous system number for the router.
- Establish an internal BGP (IBGP) session with the tunnel endpoints.
- Configure RSVP on all the routers.
- Configure OSPF or any other interior gateway protocol on all the routers.
- Configure two dynamic next-hop-based IP tunnels between the two routers.
- Configure a VRF routing instance for every router-to-host connection.

Overview

Starting in Junos OS Release 17.1, anti-spoofing capabilities are added to next-hop-based dynamic IP tunnels, where checks are implemented for the traffic coming through the tunnel to the routing instance using reverse path forwarding in the Packet Forwarding Engine.

Currently, when the gateway router receives traffic from a tunnel, only the destination address lookup is done before forwarding. With anti-spoofing protection, the gateway router does a source address lookup of the encapsulation packet IP header in the VPN to ensure that legitimate sources are injecting traffic through their designated IP tunnels. This is called the strict mode and is the default behavior of anti-spoofing protection. To pass traffic from nondesignated tunnels, the reverse path forwarding check is enabled in the loose mode. For traffic received from nonexistent sources, the reverse path forwarding check fails for both the strict and loose modes.


Anti-spoofing is supported on VRF routing instances. To enable anti-spoofing for dynamic tunnels, include the **ip-tunnel-rpf-check** statement at the **[edit routing-instances routing-instance-name routing-options forwarding-table]** hierarchy level.

Topology

[Figure 20 on page 390](#) illustrates a sample network topology enabled with anti-spoofing protection. Routers R0, R1 and R2 are each connected to hosts Host0, Host1, and Host2, respectively. Two generic routing encapsulation (GRE) next-hop-based dynamic tunnels, Tunnel 1 and Tunnel 2 – connect Router R0 with

Routers R1 and R2, respectively. The VRF routing instance is running between each router and its connected host devices.

Figure 20: Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels



Taking as an example, three packets (Packets A, B, and C) are received on Router 0 from Router R2 through the next-hop-based dynamic GRE tunnel (Tunnel 2). The source IP address of these packets are 172.17.0.2 (Packet A), 172.18.0.2 (Packet B), and 172.20.0.2 (Packet C).

The source IP address of Packets A and B belong to Host 2 and Host 1, respectively. Packet C is a nonexistent source tunnel. The designated tunnel in this example is Tunnel 2, and the nondesignated tunnel is Tunnel 1. Therefore, the packets are processed as follows:

- **Packet A**—Because the source is coming from a designated tunnel (Tunnel 2), Packet A passes the reverse path forwarding check and is processed for forwarding through Tunnel 2.
- **Packet B**—Because the source is coming from Tunnel 1, which is a nondesignated tunnel, by default, Packet B fails the reverse path forwarding check in the strict mode. If loose mode is enabled, Packet B is allowed for forwarding.
- **Packet C**—Because the source is a nonexistent tunnel source, Packet C fails the reverse path forwarding check, and the packet is not forwarded.

Configuration

CLI Quick Configuration

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

Router R0

```

set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.1/24
set interfaces ge-0/0/2 vlan-tagging
set interfaces ge-0/0/2 unit 0 vlan-id 1
set interfaces ge-0/0/2 unit 0 family inet address 172.16.0.1/16
set interfaces lo0 unit 0 family inet address 10.1.1.1/32
set routing-options router-id 10.1.1.1
set routing-options autonomous-system 100
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels T1 source-address 192.0.2.1
set routing-options dynamic-tunnels T1 gre
set routing-options dynamic-tunnels T1 destination-networks 192.0.2.0/24
set routing-options dynamic-tunnels T2 source-address 198.51.100.1
set routing-options dynamic-tunnels T2 gre
set routing-options dynamic-tunnels T2 destination-networks 198.51.100.0/24
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 10.1.1.1
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 20.1.1.1
set protocols bgp group IBGP neighbor 30.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set routing-instances VPN1 instance-type vrf
set routing-instances VPN1 interface ge-0/0/2.0
set routing-instances VPN1 route-distinguisher 100:100
set routing-instances VPN1 vrf-target target:100:1
set routing-instances VPN1 vrf-table-label
set routing-instances VPN1 routing-options forwarding-table ip-tunnel-rpf-check mode strict
set routing-instances VPN1 protocols bgp group External type external
set routing-instances VPN1 protocols bgp group External family inet unicast
set routing-instances VPN1 protocols bgp group External peer-as 200
set routing-instances VPN1 protocols bgp group External neighbor 172.16.0.1

```

Router R1


```

set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-0/0/1 vlan-tagging
set interfaces ge-0/0/1 unit 0 vlan-id 2
set interfaces ge-0/0/1 unit 0 family inet address 172.18.0.1/16
set interfaces lo0 unit 0 family inet address 20.1.1.1/32
set routing-options router-id 20.1.1.1
set routing-options autonomous-system 100
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels T1 source-address 192.0.2.2
set routing-options dynamic-tunnels T1 gre
set routing-options dynamic-tunnels T1 destination-networks 192.0.2.0/24
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 20.1.1.1
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 30.1.1.1
set protocols bgp group IBGP neighbor 10.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set routing-instances VPN2 instance-type vrf
set routing-instances VPN2 interface ge-0/0/1.0
set routing-instances VPN2 route-distinguisher 100:200
set routing-instances VPN2 vrf-target target:200:1
set routing-instances VPN2 vrf-table-label

```

R2

```

set interfaces ge-0/0/1 unit 0 family inet address 198.51.100.2/24
set interfaces ge-0/0/2 vlan-tagging
set interfaces ge-0/0/2 unit 0 vlan-id 3
set interfaces ge-0/0/2 unit 0 family inet address 172.17.0.1/16
set interfaces lo0 unit 0 family inet address 30.1.1.1/32
set routing-options router-id 30.1.1.1
set routing-options autonomous-system 100
set routing-options dynamic-tunnels gre next-hop-based-tunnel
set routing-options dynamic-tunnels T2 source-address 198.51.100.2
set routing-options dynamic-tunnels T2 gre
set routing-options dynamic-tunnels T2 destination-networks 198.51.100.0/24

```



```

set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols bgp group IBGP type internal
set protocols bgp group IBGP local-address 30.1.1.1
set protocols bgp group IBGP family inet-vpn unicast
set protocols bgp group IBGP neighbor 20.1.1.1
set protocols bgp group IBGP neighbor 10.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set routing-instances VPN3 instance-type vrf
set routing-instances VPN3 interface ge-0/0/2.0
set routing-instances VPN3 route-distinguisher 100:300
set routing-instances VPN3 vrf-target target:300:1
set routing-instances VPN3 vrf-table-label

```

Step-by-Step Procedure

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

To configure Router R0:

1. Configure Router R0's interfaces, including the loopback interface.

```

[edit interfaces]
user@R0# set ge-0/0/0 unit 0 family inet address 192.0.2.1/24
user@R0# set ge-0/0/1 unit 0 family inet address 198.51.100.1/24
user@R0# set ge-0/0/2 vlan-tagging
user@R0# set ge-0/0/2 unit 0 vlan-id 1
user@R0# set ge-0/0/2 unit 0 family inet address 172.16.0.1/16
user@R0# set lo0 unit 0 family inet address 10.1.1.1/32

```

2. Assign the router ID and autonomous system number for Router R0.

```

[edit routing-options]
user@R0# set router-id 10.1.1.1
user@R0# set autonomous-system 100

```

3. Configure IBGP peering between the routers.


```
[edit protocols]
user@R0# set bgp group IBGP type internal
user@R0# set bgp group IBGP local-address 10.1.1.1
user@R0# set bgp group IBGP family inet-vpn unicast
user@R0# set bgp group IBGP neighbor 20.1.1.1
user@R0# set bgp group IBGP neighbor 30.1.1.1
```

4. Configure OSPF on all the interfaces of Router R0, excluding the management interface.

```
[edit protocols]
user@R0# set ospf traffic-engineering
user@R0# set ospf area 0.0.0.0 interface lo0.0 passive
user@R0# set ospf area 0.0.0.0 interface all
```

5. Configure RSVP on all the interfaces of Router R0, excluding the management interface.

```
[edit protocols]
user@R0# set rsvp interface all
user@R0# set rsvp interface fxp0.0 disable
```

6. Enable next-hop-based dynamic GRE tunnel configuration on Router R0.

```
[edit routing-options]
user@R0# set dynamic-tunnels gre next-hop-based-tunnel
```

7. Configure the dynamic GRE tunnel parameters from Router R0 to Router R1.

```
[edit routing-options]
user@R0# set dynamic-tunnels T1 source-address 192.0.2.1
user@R0# set dynamic-tunnels T1 gre
user@R0# set dynamic-tunnels T1 destination-networks 192.0.2.0/24
```

8. Configure the dynamic GRE tunnel parameters from Router R0 to Router R2.

```
[edit routing-options]
user@R0# set dynamic-tunnels T2 source-address 198.51.100.1
user@R0# set dynamic-tunnels T2 gre
user@R0# set dynamic-tunnels T2 destination-networks 198.51.100.0/24
```


9. Configure a virtual routing and forwarding (VRF) routing instance on Router R0, and assign the interface connecting to Host 1 to the VRF instance.

```
[edit routing-instances]
user@R0# set VPN1 instance-type vrf
user@R0# set VPN1 route-distinguisher 100:100
user@R0# set VPN1 vrf-target target:100:1
user@R0# set VPN1 vrf-table-label
user@R0# set VPN1 interface ge-0/0/2.0
```

10. Configure an external BGP session with Host 1 for the VRF routing instance.

```
[edit routing-instances]
user@R0# set VPN1 protocols bgp group External type external
user@R0# set VPN1 protocols bgp group External family inet unicast
user@R0# set VPN1 protocols bgp group External peer-as 200
user@R0# set VPN1 protocols bgp group External neighbor 172.16.0.1
```

11. Configure anti-spoofing protection for the VRF routing instance on Router R0. This enables reverse path forwarding check for the next-hop-based dynamic tunnels, T1 and T2, on Router 0.

```
[edit routing-instances]
user@R0# set VPN1 routing-options forwarding-table ip-tunnel-rpf-check mode strict
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-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@R0# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
```



```

        address 198.51.100.1/24;
    }
}
ge-0/0/2 {
    vlan-tagging;
    unit 0 {
        vlan-id 1;
        family inet {
            address 172.16.0.1/16;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.1.1.1/32;
        }
    }
}

```

```

user@R0# show routing-options
router-id 10.1.1.1;
autonomous-system 100;
dynamic-tunnels {
    gre next-hop-based-tunnel;
    T1 {
        source-address 192.0.2.1;
        gre;
        destination-networks {
            192.0.2.0/24;
        }
    }
    T2 {
        source-address 198.51.100.1;
        gre;
        destination-networks {
            198.51.100.0/24;
        }
    }
}

```

```

user@R0# show protocols

```



```

rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  group IBGP {
    type internal;
    local-address 10.1.1.1;
    family inet-vpn {
      unicast;
    }
    neighbor 20.1.1.1;
    neighbor 30.1.1.1;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface all;
  }
}

```

user@R0# **show routing-instances**

```

VPN1 {
  instance-type vrf;
  interface ge-0/0/2.0;
  route-distinguisher 100:100;
  vrf-target target:100:1;
  vrf-table-label;
  routing-options {
    forwarding-table {
      ip-tunnel-rpf-check {
        mode strict;
      }
    }
  }
  protocols {
    bgp {
      group External {

```



```

    type external;
    family inet {
        unicast;
    }
    peer-as 200;
    neighbor 172.16.0.1;
}
}
}
}
```

Verification

IN THIS SECTION

- [Verifying Basic Configuration | 398](#)
- [Verifying Dynamic Tunnel Configuration | 399](#)
- [Verifying Anti-Spoofing Protection Configuration | 400](#)

Confirm that the configuration is working properly.

Verifying Basic Configuration

Purpose

Verify the OSPF and BGP peering status between the Router R0 and Routers R1 and R2.

Action

From operational mode, run the **show ospf neighbor** and **show bgp summary** commands.

```
user@R0> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
192.0.2.2	ge-0/0/0.0	Full	20.1.1.1	128	32
198.51.100.2	ge-0/0/1.0	Full	30.1.1.1	128	32

```
user@R0> show bgp summary
```



```

Groups: 2 Peers: 3 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l3vpn.0
                0          0          0          0          0          0
Peer           AS        InPkt    OutPkt    OutQ     Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
20.1.1.1       100        182      178        0         0    1:20:27 Establ

    bgp.l3vpn.0: 0/0/0/0
30.1.1.1       100        230      225        0         0    1:41:51 Establ

    bgp.l3vpn.0: 0/0/0/0
172.16.0.1     200         0         0         0         0    1:42:08 Establ

```

Meaning

The OSPF and BGP sessions are up and running between the Routers R0, R1, and R2.

Verifying Dynamic Tunnel Configuration

Purpose

Verify the status of the next-hop-based dynamic GRE tunnels between the Router R0 and Routers R1 and R2.

Action

From operational mode, run the **show route table inet.3**, and the **show dynamic-tunnels database terse** commands.

```
user@R0> show route table inet.3
```

```

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24      *[Tunnel/300] 01:47:57
                  Tunnel
192.0.2.2/24      *[Tunnel/300] 01:47:57
                  Tunnel Composite
198.51.100.0/24  *[Tunnel/300] 01:47:57
                  Tunnel

```



```
198.51.100.2/24      *[Tunnel/300] 01:47:57
                    Tunnel Composite
```

user@R0> **show dynamic-tunnels database terse**

```
Table: inet.3

Destination-network: 192.0.2.0/24
Destination      Source      Next-hop      Type      Status
192.0.2.2/24     192.0.2.1   0xb395e70     nhid 612   gre      Up

Destination-network: 198.51.100.0/24
Destination      Source      Next-hop      Type      Status
198.51.100.2     198.51.100.1 0xb395e70     nhid 612   gre      Up
```

Meaning

The two next-hop-based dynamic GRE tunnels, Tunnel 1 and Tunnel 2, are up.

Verifying Anti-Spoofing Protection Configuration

Purpose

Verify that the reverse path forwarding check has been enabled on the VRF routing instance on Router R0.

Action

From the operational mode, run the **show krt table VPN1.inet.0 detail**.

user@R0> **show krt table VPN1.inet.0 detail**

```
KRT tables:
VPN1.inet.0          : GF: 1 krt-index: 8      ID: 0 kernel-id: 8
  flags: (null)
  tunnel rpf config data : enable, strict, filter [0], 0x2
  tunnel rpf tlv data : enable, strict, filter [0], 0x4
  unicast reverse path: disabled
  fast-reroute-priority: 0
  Permanent NextHops
    Multicast          : 0 Broadcast : 0
    Receive            : 0 Discard   : 0
    Multicast Discard: 0 Reject     : 0
```



```

Local          : 0    Deny          : 0
Table          : 0

```

Meaning

The configured reverse path forwarding check is enabled on the VRF routing instance in the strict mode.

RELATED DOCUMENTATION

[Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview | 385](#)
[ip-tunnel-rpf-check | 2632](#)

Next-Hop-Based Dynamic Tunnel Localization Overview

IN THIS SECTION

- [Benefits of Next-Hop-Based Dynamic Tunnel Localization | 402](#)
- [Use Cases for Next-Hop-Based Dynamic Tunnel Localization | 402](#)
- [Traffic Handling with Localization of Next-Hop-Based Dynamic Tunnels | 402](#)
- [Configuring Next-Hop-Based Dynamic Tunnels Localization | 403](#)
- [Troubleshooting Localized Next-Hop-Based Dynamic Tunnels | 406](#)
- [Unsupported Features for Next-Hop-Based Dynamic Tunnels Localization | 407](#)

Next-hop-based dynamic tunnels include generic routing encapsulation (GRE) tunnels and MPLS-over-UDP tunnels. These tunnels provide a scaling advantage over the interface-based tunnels. However, unlike the interface-based tunnels, the next-hop-based dynamic tunnels are anchorless in nature, where the forwarding information of the tunnels is distributed to the Packet Forwarding Engines (PFEs) on every line card on the device. This limits the maximum number of tunnels supported on the device to the tunnel capacity of a single line card. With the support for localization, you can configure next-hop-based dynamic tunnel localization to create the forwarding information only on the PFE of a line card that is designated as the anchor PFE. The PFEs on the other line cards on the device have state forwarding information to steer the packets to the anchor PFE. This provides a scaling advantage by increasing the maximum number of tunnels supported on a device.

Benefits of Next-Hop-Based Dynamic Tunnel Localization

Provides a scaling advantage by increasing the maximum number of tunnels supported on a device.

Use Cases for Next-Hop-Based Dynamic Tunnel Localization

- The IPsec gateway devices that host a number of MS-MPC are used to terminate IPsec tunnels and are required to support moderate load. This support is affected with the use of next-hop-based dynamic tunnels when the scaling limit of the device is reached. With the localization of next-hop-based dynamic tunnels, the maximum number of the tunnels supported is increased, allowing the device to accommodate more tunnels at the cost of an extra fabric hop.
- For Internet or VPN gateway devices, such as a virtual public cloud data center, there is a need for the gateway devices to communicate with a large number of servers. The data center servers are reachable through next-hop-based dynamic tunnels. The anchorless property of the dynamic tunnels limits the overall scaling numbers of the device. The gateway devices host multiple MPCs, with increased traffic demands. With the localization of the next-hop-based dynamic tunnels, the tunnels can be spread across the MPCs, thereby facilitating an increase in the tunnel scaling numbers.

Traffic Handling with Localization of Next-Hop-Based Dynamic Tunnels

With support for localization, the next-hop-based dynamic tunnel state is localized to an anchor Packet Forwarding Engine, and the other Packet Forwarding Engine has the tunnel state for steering traffic to the tunnel anchor.

Figure 21 on page 402 illustrates the forwarding path of next-hop-based dynamic tunnels without localization.

Figure 21: Forwarding Path of Next-Hop-Based Dynamic Tunnels Without Localization


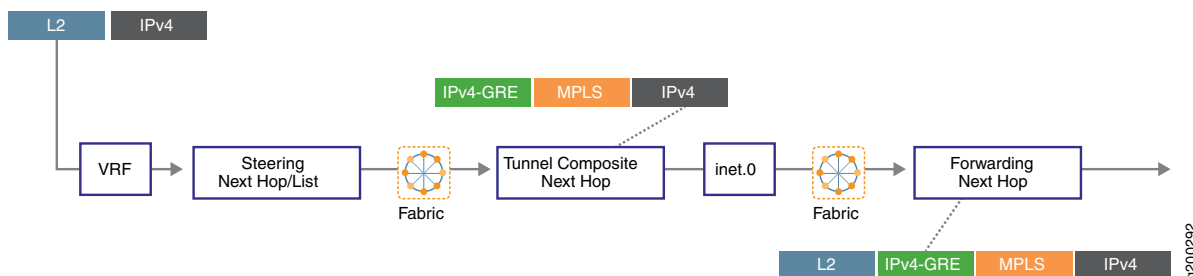


Figure 22 on page 403 illustrates the forwarding path of next-hop-based dynamic tunnels with localization.

Figure 22: Forwarding Path of Next-Hop-Based Dynamic Tunnels With Localization



Configuring Next-Hop-Based Dynamic Tunnels Localization

IN THIS SECTION

- [Configuring Localization for New Next-Hop-Based Dynamic Tunnels | 403](#)
- [Configuring Localization for Existing Next-Hop-Based Dynamic Tunnels | 405](#)

Localization support can be configured for newly created next-hop-based dynamic tunnels, or for existing non-local dynamic tunnels.

Configuring Localization for New Next-Hop-Based Dynamic Tunnels

The localization of next-hop-based dynamic tunnels uses a policy-based approach to specify prefix groups. In other words, route policies are used to apply the localization properties to the next-hop-based dynamic tunnels. Dynamic tunnel attribute profiles are created and configured under routing options for association with the prefix group using the policy.

1. Creating dynamic tunnel profiles.

The dynamic tunnel profile specifies the tunnel type and the anchor Packet Forwarding Engine information. Multiple dynamic tunnel profiles can be created for localization of the dynamic tunnels.

The values for the dynamic tunnel type can be GRE, UDP, or BGP-SIGNAL.

Although BGP-SIGNAL is not a valid tunnel type, on assigning BGP-SIGNAL as the tunnel type, the tunnels created from the BGP-signalled attributes are localized. When using BGP-SIGNAL, the tunnel type is decided based on the type advertised by BGP in its TLV. BGP-SIGNAL tunnels are always next-hop-based tunnels. The GRE tunnels created dynamically by BGP-SIGNAL are always next-hop-based, even if the user has manually configured tunnels created by GRE to use IFLs.

The anchor Packet Forwarding Engine value is the line card of the anchor Packet Forwarding Engine, for example, `pfe-x/y/0`. This information can be viewed from the **show interfaces terse pfe*** command output.

Sample Configuration:

```
[edit routing-options]
dynamic-tunnels {
  dynamic-tunnel-attributes attribute-1 {
    dynamic-tunnel-type <GRE | UDP | BGP-SIGNAL>;
    dynamic-tunnel-anchor-pfe pfe-1/0/0;
  }
}
```

2. Associating dynamic tunnel profile to prefix list.

Configuring a policy with **dynamic-tunnel-attributes** as the action associates the dynamic tunnel to the prefix list. The policy **from** action allows the creation of tunnel with specified attributes for any matching condition, such as a prefix range, community, or source address of BGP routes, and so on.

Sample configuration:

```
[edit policy-options]
policy-statement policy-name {
  term term {
    from {
      <route-filter | next-hop | community>>;
    }
    then {
      dynamic-tunnel-attributes <attribute-name>;
    }
  }
}
```

3. Including the tunnel policy under the forwarding table export policy.

After the policy is configured, it is included in the forwarding table export policy for the parsing of the policy.

Using the export-policy, the tunnel attributes get associated with the route. Whenever a route from BGP is queued for resolution, the forwarding table export policy is evaluated, and the tunnel attributes are obtained from the policy module based on the applied filters. The obtained tunnel attributes are then attached to the next hop in form of a tunnel composite next hop. The corresponding anchor forwarding structures, based on the Packet Forwarding Engine name and tunnel type, are created and sent to the forwarding table before a tunnel composite next hop is sent. However, if none of the

attributes map to the tunnel composite next hop, then the forwarding structure is created on every Packet Forwarding Engine, similar to the non-localized dynamic tunnels.

Sample configuration:

```
[edit routing-options]
forwarding-table {
  export dynamic-tunnel;
}
```

Configuring Localization for Existing Next-Hop-Based Dynamic Tunnels



CAUTION: Making on the fly changes to dynamic tunnel attributes can result in an FPC crash due to high memory utilization. Hence, we recommend deactivating the dynamic-tunnels configuration before configuring localization.

To update tunnel attributes for existing next-hop-based dynamic tunnels, the following should be performed:

1. Deactivate **dynamic-tunnels** configuration under the **[edit routing-options]** hierarchy level.

Sample configuration:

```
[edit routing-options]
user@host# deactivate dynamic-tunnels
user@host# commit
```

2. Change tunnel attributes as required.
3. Activate **dynamic-tunnels** configuration under the **[edit routing-options]** hierarchy level.

Sample configuration:

```
[edit routing-options]
user@host# activate dynamic-tunnels
user@host# commit
```


To configure localization for existing non-local next-hop-based dynamic tunnels:



CAUTION: Making on the fly changes to configure localization for existing non-local next-hop-based dynamic tunnels can result in an FPC crash due to high memory utilization. Hence, we recommend deactivating the dynamic-tunnels configuration before configuring localization.

1. Deactivate the **dynamic-tunnels** configuration at the **[edit routing-options]** hierarchy level.
2. Create tunnel-attributes profile and add policy for localizing the dynamic tunnels, similar to new next-hop-based dynamic tunnels.
3. Activate the **dynamic-tunnels** configuration.

Troubleshooting Localized Next-Hop-Based Dynamic Tunnels

With localization of next-hop-based dynamic tunnels, the tunnel composite next hops are associated with anchor Packet Forwarding Engine IDs. The following traceroute configuration statements at the **[edit routing-options]** hierarchy level help in troubleshooting the localized dynamic tunnels:

- **dynamic-tunnels traceoptions flag all**—Tracking creation and deletion of tunnel in DTM.
- **resolution traceoptions flag tunnel**—Tracking resolver operations on BGP route.
- **forwarding-table traceoptions flag all**—Tracking tunnels sent to the kernel.
- **traceoptions flag all**—Tracking of route learning process.

The following commands can be used to check if a route is using a localized next-hop-based dynamic tunnel:

1. **show route *prefix* extensive**—To obtain the indirect next hop.

For example:

```
user@host> show route 1.2.3.4 extensive
```

```
MPLS-over-UDP-PE1.inet.0: 24 destinations, 26 routes (24 active, 0 holddown, 0
  hidden)
1.2.3.4/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.2.3.4/32 -> {indirect(1048577)}
Page 0 idx 1, (group pe1-cel type External) Type 1 val 0xb209a78 (adv_entry)
  Advertised metrics:
```



```

Nexthop: Self
AS path: [100] I
Communities: target:600:1 encapsulation:mpls-in-udp(0xd)

```

2. **show krt indirect-next-hop index *indirect-next-hop* detail**—To check for anchor Packet Forwarding Engine field in the detailed output of the indirect next hop.

For example:

```
user@host> show krt indirect-next-hop index 1048577 detail
```

```

Indirect Nexthop detail:
Index: 1048577 Protocol next-hop address: 1.1.1.6
  RIB Table: bgp.l3vpn.0          Label: Push 299808
  Policy Version: 2              References: 11
  Locks: 3                      0xb227980
  Flags: 0x0
  INH Session ID: 0x0
  Ref RIB Table: unknown
  Export policy detail:
  (Dynamic tunnel hash : 309985522)
    Tunnel type: UDP, Reference count: 4, nhid: 1016
    Destination address: 1.1.1.6, Source address: 1.1.1.2
    Anchored-PFE: pfe-1/0/0
    VPN Label: Push 299808, TTL action: prop-ttl
  IGP FRR Interesting proto count : 11
  Chain IGP FRR Node Num          : 1
    IGP Resolver node(hex)        : 0xc838b94
    IGP Route handle(hex)         : 0xb1d7674   IGP rt_entry protocol      : Tunnel

    IGP Actual Route handle(hex) : 0x0          IGP Actual rt_entry protocol
: Any

```

Unsupported Features for Next-Hop-Based Dynamic Tunnels Localization

Junos OS does not support the following functionality with localization for next-hop-based dynamic tunnels:

- Chained composite next hops at the **[edit routing-options forwarding-table chained-composite-next-hop ingress l3vpn]** hierarchy level.
- Anchor Packet Forwarding Engine resiliency.

There is no resiliency support for next-hop-based dynamic tunnels with localization. After localization of the next-hop-based dynamic tunnels, the anchor Packet Forwarding Engine becomes the single entity

for processing any given tunnel on the device. Although anchor Packer Forwarding Engine resiliency is not supported, for gateway devices, redundancy at the gateway device ensures that when the Packer Forwarding Engine to which the tunnel composite next hop is delegated goes down, the traffic must be rerouted to the redundant gateway device. The routing protocol process monitors the state of the Packer Forwarding Engine, and withdraws BGP advertisement of all the routes pointing to the tunnel composite next hops anchored on that Packer Forwarding Engine.

Only the anchored Packet Forwarding Engine has the full-fledged tunnel composite next hop and all the other Packet Forwarding Engines have only steering entries to forward traffic to the anchor Packet Forwarding Engine. These steering entries are not withdrawn, when an anchor FPC goes down.

- Localization of next-hop-based dynamic tunnels is not supported on logical systems.
- IPv6 is not supported with localization of next-hop-based dynamic tunnels.
- With localization, the **show dynamic-tunnels database summary** command does not display accurate tunnels summary when the state of the anchor Packet Forwarding Engine line card is not up. As a workaround, use the **show dynamic-tunnels database** and **show dynamic-tunnels database terse** command output.

RELATED DOCUMENTATION

[Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels](#) | 367

Example: Configuring a Next-Hop-Based Dynamic GRE Tunnels

Configuring Poloccy-Based Multipath Routes

IN THIS CHAPTER

- [Policy-Based Multipath Routes Overview | 409](#)

Policy-Based Multipath Routes Overview

IN THIS SECTION

- [Understanding Policy-Based Multipath Routes | 409](#)
- [Benefits of Policy-Based Multipath Routes | 410](#)
- [Policy-Based Multipath Routes for Route Resolution | 410](#)
- [Sample Route Resolution Using Policy-Based Multipath Routes | 411](#)
- [Enhancement to Class-of-Service \(CoS\) Forwarding-Policy | 413](#)
- [Enhancements to Policy Match Protocol | 414](#)
- [Impact of Configuring Policy-Based Multipath Route on Network Performance | 415](#)

Segment routing networks can have multiple transport protocols in the core. You can combine segment routing SR-TE LDP or RSVP routes and SR-TE IP routes and install a multipath route in the routing information base (also known as routing table). You can then steer selective service traffic using the mutlipath route through policy configuration.

Understanding Policy-Based Multipath Routes

There are different transport protocols in a network, such as IGP, labelled IGP, RSVP, LDP, and segment routing traffic-engineering (SR-TE) protocols, that are used to resolve service traffic. However, you could not use a combination of the transport protocols to resolve the service traffic. With the introduction of the policy-based multipath feature, you can combine segment routing traffic-engineered (SR-TE) LDP or RSVP routes and SR-TE IP routes to create a multipath route that is installed in the routing information

base. You can resolve BGP service routes over the multipath route through policy configuration and steer traffic differently for different prefixes.

A multipath route has combined next hops of route entries that are used for load balancing. All the supporting routes of the multipath route entry must be in same routing information base. When the supporting routes are under different routing information base, you can use the **rib-group** configuration statement to add route entries to a particular routing information base.

You can configure a multipath route using a policy to select the list of routes whose next hops is to be combined together. When you include the **policy-multipath** statement along with the **policy** statement at the **[edit routing-options rib routing-table-name]** hierarchy level, a policy-based multipath route is created.

The policy-based multipath feature is supported for both IP and IPv6 protocols, and can be configured under the **[edit routing-instances]** hierarchy level.

For example:

```
[edit routing-options]
user@host# set rib inet.3 policy-multipath policy example-policy
[edit policy-options]
user@host# set policy-statement example-policy from example-conditions
user@host# set policy-options policy-statement example-policy then accept
```

The configured policy is applied to each route entry for a given prefix. The multipath route is created only when more than one route (including active route) passes the policy. Any action commands configured in the policy, such as **apply**, is evaluated using the active route. For non-active routes, the policy is applied to check if the routes can participate in the multipath route or not. Multipath routes inherit all attributes of the active route. These attributes can be modified using the multipath policy configuration.

Benefits of Policy-Based Multipath Routes

- Provides flexibility to combine core network protocols to steer selective traffic.
- Optimizes network performance with weighted equal-cost multipath using multipath routes.

Policy-Based Multipath Routes for Route Resolution

You can combine segment routing traffic-engineered (SR-TE) LDP or RSVP routes and SR-TE IP routes and install a multipath route in the routing information base. The policy-based multipath routes are not active entries in the routing information base. When a multipath route is generated by configuration of policy, it is used for resolving protocol next hops instead of active routes. A multipath route next hop is created by merging gateways of next hops of each constituent route.

Take the following into consideration when configuring policy-based multipath routes for route resolution:

- If the member route of a multipath route points to a next hop other than the router next hop or an indirect next hop with forwarding next hop to the router next hop, such next hops are ignored.
- If the constituent routes point to indirect next hop, then gateways from the forwarding-next hop are merged and the indirect next hop is ignored.
- If total number of gateways exceeds the **maximum-ecmp** supported on the device, then only the **maximum-ecmp** gateways are retained and all other gateways are ignored.
- Gateways with lower weights are given preference. When one of the member route has unilist of indirect next hops and each of the next hop is pointing to a forwarding next hop, there can be weight values both at the indirect next hop and at forwarding next hop. In such cases, weight value of gateways is updated to reflect the combined effect of weights at both levels.

Sample Route Resolution Using Policy-Based Multipath Routes

Taking as an example, let us assume there are segment routing traffic-engineered LSPs, label IS-IS routes, and LDP LSPs for a destination 1.1.1.1/32, as displayed in the output below:

```
1.1.1.1/32      *[SPRING-TE/8] 00:00:58, metric 1, metric2 30
                 > to 13.1.1.2 via ge-0/0/1.1, Push 33333, Push 801005, Push
801006(top)

                 [L-ISIS/14] 1w0d 00:15:57, metric 10
                 > to 12.1.1.1 via ge-0/0/0.1
                   to 12.22.1.1 via ge-0/0/0.2
                   to 12.23.1.1 via ge-0/0/0.3
                   to 12.24.1.1 via ge-0/0/0.4
                   to 12.25.1.1 via ge-0/0/0.5
                   to 13.1.1.2 via ge-0/0/1.1, Push 801001, Push 801005(top)
                 [LDP/19] 1w0d 00:09:27, metric 1
                 > to 12.1.1.1 via ge-0/0/0.1
                   to 12.22.1.1 via ge-0/0/0.2
                   to 12.23.1.1 via ge-0/0/0.3
                   to 12.24.1.1 via ge-0/0/0.4
                   to 12.25.1.1 via ge-0/0/0.5
                   to 13.1.1.2 via ge-0/0/1.1, Push 801001, Push 801005(top)
```

Here, segment routing LSP is the active route entry to the 1.1.1.1 destination, and by default, only this route is used to resolve any services resolving over 1.1.1.1.

When there is a requirement to use more than one protocols for resolving service routes, you can achieve this by configuring **policy-multipath** to combine the protocols. For instance, if segment routing and LDP

paths are required for service resolution, you must configure **policy-multipath** combining the segment routing and LDP routes for prefix 1.1.1.1.

For example:

```
[edit policy-options]
user@host# set rib inet.3 policy-multipath policy example-policy
user@host# set policy-statement abc term 1 from protocol spring-te
user@host# set policy-statement abc term 1 from protocol ldp
user@host# set policy-statement abc term 1 from route-filter 1.1.1.1/32 exact
user@host# set policy-statement abc term 1 then accept
```

With this configuration, you create a policy-based multipath route for prefix 1.1.1.1/32 that uses constituent route entries of segment routing and LDP protocols.

You can view the multipath route using the **show route** command output, as follows:

```
1.1.1.1/32          *[SPRING-TE/8] 00:10:28, metric 1, metric2 30
> to 13.1.1.2 via ge-0/0/1.1, Push 33333, Push 801005, Push
801006(top)

[L-ISIS/14] 1w0d 00:25:27, metric 10
> to 12.1.1.1 via ge-0/0/0.1
  to 12.22.1.1 via ge-0/0/0.2
  to 12.23.1.1 via ge-0/0/0.3
  to 12.24.1.1 via ge-0/0/0.4
  to 12.25.1.1 via ge-0/0/0.5
  to 13.1.1.2 via ge-0/0/1.1, Push 801001, Push 801005(top)
[LDP/19] 1w0d 00:18:57, metric 1
> to 12.1.1.1 via ge-0/0/0.1
  to 12.22.1.1 via ge-0/0/0.2
  to 12.23.1.1 via ge-0/0/0.3
  to 12.24.1.1 via ge-0/0/0.4
  to 12.25.1.1 via ge-0/0/0.5
  to 13.1.1.2 via ge-0/0/1.1, Push 801001, Push 801005(top)
[Multipath/8] 00:03:13, metric 1, metric2 30
> to 12.1.1.1 via ge-0/0/0.1
  to 12.22.1.1 via ge-0/0/0.2
  to 12.23.1.1 via ge-0/0/0.3
  to 12.24.1.1 via ge-0/0/0.4
  to 12.25.1.1 via ge-0/0/0.5
  to 13.1.1.2 via ge-0/0/1.1, Push 33333, Push 801005, Push
801006(top)

  to 13.1.1.2 via ge-0/0/1.1, Push 801001, Push 801005(top)
```


You can see from the command output that the multipath route combines next hops of segment routing and LDP paths. The multipath route is not active, and by default, the route preference and metric is the same as that of active route.

NOTE:

You can use the following combinations for the policy-based multipath route: However we cannot create multipath of LDP/L-ISIS as active-route is not part of multipath.

- Segment routing traffic-engineered LSPs and LDP LSPs.
- Segment routing traffic-engineered LSPs, and label IS-IS paths.
- Segment routing traffic-engineered LSPs, LDP LSPs, and label IS-IS paths.

However, you cannot create multipath route of LDP and label IS-IS, as the active route is not part of the multipath route.

With the same configuration, assuming that there is a static route 1.2.3.4/32 configured with a protocol next hop of 1.1.1.1, this route is resolved using the multipath route over both segment routing traffic-engineered LSPs and LDP LSPs.

For example:

```
1.2.3.4/32          *[Static/5] 00:00:12, metric2 1
                    to 12.1.1.1 via ge-0/0/0.1
                    > to 12.22.1.1 via ge-0/0/0.2
                    to 12.23.1.1 via ge-0/0/0.3
                    to 12.24.1.1 via ge-0/0/0.4
                    to 12.25.1.1 via ge-0/0/0.5
                    to 13.1.1.2 via ge-0/0/1.1, Push 33333, Push 801005, Push
801006(top)
                    to 13.1.1.2 via ge-0/0/1.1, Push 801001, Push 801005(top)
```

Enhancement to Class-of-Service (CoS) Forwarding-Policy

For class-of-service-based forwarding, you must use the **forwarding-policy next-hop-map** configuration statement.

Prior to Junos OS Release 19.1R1, the match conditions supported under class-of-service-based forwarding included:

- **next-hop**—Match next hop based on outgoing interface or next hop address.

- **lsp-next-hop**—Match named LSPs using regular expression of LSP name.
- **non-lsp-next-hop**—Match all LSPs without an LSP name.

With the policy-based multipath route feature, you can also match all next hops without a label for certain prefixes. To do this, you must enable the **non-labelled-next-hop** option at the **[edit class-of-service forwarding-policy next-hop-map *map-name* forwarding-class *forwarding-class-name*** hierarchy level.

For example:

```
[edit]
class-of-service {
  forwarding-policy {
    next-hop-map abc {
      forwarding-class best-effort {
        non-labelled-next-hop;
      }
    }
  }
}
```

Enhancements to Policy Match Protocol

Prior to Junos OS Release 19.1R1, when you used a policy to match protocol using the **from protocol** statement at the **[edit policy-options policy-statement *statement-name*]** hierarchy level, all protocol routes (labeled and unlabeled) were matched. With the policy-based multipath route feature, you can match labeled protocol routes specifically.

The options for matching labeled protocols are:

- **l-isis**—Match labeled IS-IS routes. The **isis** option matches IS-IS routes, excluding label IS-IS routes.
- **l-ospf**—Match labeled OSPF routes. The **ospf** option matches all OSPF routes, including OSPFv2, OSPFv3 and label OSPF.

For example:

```
[edit]
policy-options {
  policy-statement abc {
    from protocol [ l-ospf l-isis ];
  }
}
```


Impact of Configuring Policy-Based Multipath Route on Network Performance

When you configure policy-based multipath route, a change of route in the routing information base results in the evaluation of the policy to check if a multipath route needs to be created. Because this feature requires that member routes must be in the same routing information base, the **rib-group** statement is used to merge routes from different routing information base. Configuring the **rib-group** statement at the application level increases number of routes in the system.

When there are a number of routes in the routing information base, constant change of routes leads to reevaluation of the multipath policy. This could impact network performance. It is recommended to configure the policy-based multipath route feature only when required.

RELATED DOCUMENTATION

| [policy-multipath](#) | 2731

3

PART

MPLS Label-Switched Paths

MPLS Label Operations | **419**

MPLS LSP Routes | **473**

MPLS LSP Routers | **513**

Configuring MPLS LSPs | **525**

Configuring Point-to-Multipoint LSPs | **691**

Configuring Container LSPs | **729**

Configuring Pop-and-Forward LSPs | **797**

MPLS Label Operations

IN THIS CHAPTER

- [MPLS Label Overview | 419](#)
- [MPLS Label Allocation | 420](#)
- [Operations on MPLS Labels | 421](#)
- [Understanding MPLS Label Operations | 422](#)
- [Understanding MPLS Label Manager | 425](#)
- [Understanding MPLS Label Operations on EX Series Switches | 426](#)
- [How a Packet Travels Along an LSP | 430](#)
- [Types of LSPs | 431](#)
- [Scope of LSPs | 431](#)
- [Special MPLS Labels | 432](#)
- [Entropy Label Support in Mixed Mode Overview | 433](#)
- [Abstract Hops for MPLS LSPs Overview | 433](#)
- [Example: Configuring Abstract Hops for MPLS LSPs | 446](#)
- [Configuring the Maximum Number of MPLS Labels | 467](#)
- [Configuring MPLS to Pop the Label on the Ultimate-Hop Router | 470](#)
- [Advertising Explicit Null Labels to BGP Peers | 471](#)

MPLS Label Overview

Packets traveling along an LSP are identified by a label—a 20-bit, unsigned integer in the range 0 through 1,048,575. For push labels on ingress routers, no labels in this range are restricted. For incoming labels on the transit static LSP, the label value is restricted to 1,000,000 through 1,048,575.

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

MPLS Label Allocation


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

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

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


[Figure 23 on page 420](#) shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 23: Label Encoding



[Figure 24 on page 421](#) illustrates the purpose of the class-of-service bits (also known as the EXP or experimental bits). Bits 20 and 21 specify the queue number. Bit 22 is the packet loss priority (PLP) bit used to specify the random early detection (RED) drop profile. For more information about class of service and the class-of-service bits, see [“Configuring Class of Service for MPLS LSPs” on page 1311](#).

Figure 24: Class-of-Service Bits



RELATED DOCUMENTATION

| [per-prefix-label](#) | 2725

Operations on MPLS Labels

The router supports the following label operations:

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

Understanding MPLS Label Operations

IN THIS SECTION

- [MPLS Label-Switched Paths and MPLS Labels | 422](#)
- [Reserved Labels | 423](#)
- [MPLS Label Operations | 423](#)
- [Penultimate-Hop Popping and Ultimate-Hop Popping | 425](#)

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

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 23 on page 420](#) shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 25: 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 26 on page 424 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 26: MPLS Label Swapping




Figure 26 on page 424 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 | 38](#)

[Configuring MPLS on Provider Edge Switches | 83](#)

[Configuring MPLS on Provider Switches | 88](#)

MPLS Applications User Guide

Junos OS VPNs Library for Routing Devices

Understanding MPLS Label Manager

MPLS label manager is used to manage different label types such as LSI, dynamic, block, and static, which are supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. These line cards provide more flexibility and scalability, when the **enhanced-ip** command is configured on the device.

The existing behavior of **label-space** command is retained, which is *not recommended*. To provide additional functionality such as multiple ranges for each type of label, **label-range** command is introduced under the **[edit protocols mpls label usage]** hierarchy, which is independent of **label-space** configuration. You can choose either style if only one range is needed for each type of label.

The following features are optimized with the **enhanced-ip** command configured on the device:

- Allows you to define the system wide global label pool to be used by segment-routing global block (SRGB) through IS-IS routing protocol.
- Increases the **vrf-table-label** space to at least 16,000, if the platform can support the scale.
- Allows you to specify the label value to be used by static VRF table label.
- Allows you to specify the label value range to be used by supported label application types.
- Allows you to change dynamically the SRGB and label type ranges.

RELATED DOCUMENTATION

[vrf-table-label](#) | 2826

static

[show mpls label usage](#) | 3207

[show mpls label usage label-range](#) | 3211

Understanding MPLS Label Operations on EX Series Switches

IN THIS SECTION

- [MPLS Label-Switched Paths and MPLS Labels on the Switches](#) | 427
- [Reserved Labels](#) | 428
- [MPLS Label Operations on the Switches](#) | 428
- [Penultimate-Hop Popping and Ultimate-Hop Popping](#) | 429

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 again forwarded based on its IP routing information.

This topic describes:


MPLS Label-Switched Paths and MPLS Labels on the Switches

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.

NOTE: The implementation of MPLS on Juniper Networks EX3200 and EX4200 Ethernet Switches supports only single-label packets. However, MPLS on Juniper Networks EX8200 Ethernet Switches supports packets with as many as three labels.

Figure 23 on page 420 shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 27: 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 the switches:

- 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.
- 2, IPv6 Explicit Null label—This value is legal only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt.
- 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 on the Switches

EX Series switches support the following label operations:

- Push
- Pop

- Swap

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 26 on page 424 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 to 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 28: MPLS Label Swapping




Figure 26 on page 424 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.

On EX8200 switches, you can choose to use either the default, PHP, or to configure ultimate-hop popping.

- 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

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

Junos OS VPNs Library for Routing Devices

MPLS Applications User Guide

How a Packet Travels Along an LSP

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

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

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

Types of LSPs

There are three types of LSPs:

- Static LSPs—For static paths, you must manually assign labels on all routers involved (ingress, transit, and egress). No signaling protocol is needed. This procedure is similar to configuring static routes on individual routers. Like static routes, there is no error reporting, liveliness detection, or statistics reporting.
- LDP-signaled LSPs—See [“LDP Introduction” on page 918](#).
- RSVP-signaled LSPs—For signaled paths, RSVP is used to set up the path and dynamically assign labels. (RSVP signaling messages are used to set up signaled paths.) You configure only the ingress router. The transit and egress routers accept signaling information from the ingress router, and they set up and maintain the LSP cooperatively. Any errors encountered while establishing an LSP are reported to the ingress router for diagnostics. For signaled LSPs to work, a version of RSVP that supports tunnel extensions must be enabled on all routers.

There are two types of RSVP-signaled LSPs:

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

Scope of LSPs

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

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

Special MPLS Labels

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

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

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

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

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

Entropy Label Support in Mixed Mode Overview

Starting with Junos OS Release 14.2, entropy label is supported in mixed mode chassis where the entropy label can be configured without enhanced-ip configuration. The entropy label helps transit routers load-balance MPLS traffic across ECMP paths or link aggregation groups. The entropy label introduces a load-balancing label to be used by routers to load balance traffic rather than relying on deep packet inspection, reducing the packet processing requirements in the forwarding plane at the expense of increased label stack depth. Junos OS supports the entropy label only for MX Series routers with MPCs or MICs and can be enabled with enhanced-ip mode. But, this leads to a packet drop if the core-facing interface has an entropy label configured on the MPC or MIC and the other end of this core-facing connection has a DPC line card. In order to avoid this, the entropy label is now supported in mixed mode where the entropy label can be configured without enhanced-ip configuration. This allows MX Series router DPCs to support a pop out entropy label. However, this does not support a flow label.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, entropy label is supported in mixed mode chassis where the entropy label can be configured without enhanced-ip configuration.

Abstract Hops for MPLS LSPs Overview

IN THIS SECTION

- [Understanding Abstract Hops | 434](#)
- [Benefits of Using Abstract Hops | 435](#)
- [Junos OS Implementation of Abstract Hops | 437](#)

An abstract hop is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs), which results in a user-defined group or cluster of routers that can be sequenced and used as constraints for setting up an MPLS label-switched path (LSP). Abstract hops overcome the limitations of existing path constraint specifications and provide several benefits to the traffic engineering capabilities of MPLS.

Understanding Abstract Hops

The path constraint for setting up of an MPLS LSP can be specified as either individual routers in the form of real hops or as a set of routers by way of administrative group or color specification. When a path constraint uses real hops (strict or loose), the LSP is set up along a specified sequence of routers (for example, R1, R2, ... Rn). When a path constraint uses an administrative group or color specification, a group of routers that meet the specified criteria is used to set up the LSP without picking a specific router, and unlike real-hop constraint, there is no sequence among the different groups of routers used in the constraint.

The drawback of real-hop constraint is that in a failure scenario, if any of the router hops goes down or the bandwidth utilization of the attached interface gets saturated, the path goes down (or relies on local or end-to-end protection). Although other alternative routers might be available to recover or set up the LSP, the LSP remains down until the operator configures another router hop sequence as the path constraint to bring the path up again or to disengage the protection path.

The administrative group or color specification constraint overcomes this limitation of a real-hop constraint to a certain extent. Here, when one of the routers in the group goes down or has its link capacity saturated, setting up of the LSP is not affected. This is because the next hop router to be used in the path constraint is not picked beforehand, and the LSP is set up along other routers that have the same administrative group or color without operator intervention. However, the drawback with router group constraints is that a sequence cannot be specified among the hop constraints.

Abstract hops overcome these drawbacks by creating user-defined router groups, where each member router meets a user-defined constraint. The user-defined constraint is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs). Ordering is achieved among the router groups by specifying a sequence of abstract hops used in a path constraint. As a result, abstract hops combine the ordering property of real-hop constraint specification and the resilience that comes with the other traffic engineering constraints.

A path can use a combination of real and abstract hops as constraints. When using abstract hops, instead of specifying a sequence of routers (R1, R2, ... Rn) as with real hops, you specify an ordered set of router groups or abstract hops (G1, G2, ... Gn) as the path constraint. Each specified router group, Gi for example, consists of some user-defined set of routers—R1, R2, Rj, ... Rn. When one of the routers in the group goes down, say Router Rj in group Gi, another router, say Router Rk, from the same group Gi is picked up by path computation to replace the router that went down (that is, Router Rj). This is because the path constraint is sequenced and has to go through a sequence of abstract hops, instead of a sequence of individual routers.

Benefits of Using Abstract Hops

IN THIS SECTION

- [Specifying a Sequence of Constraint Combinations | 435](#)
- [Avoiding New Network Configuration on Transit Nodes | 435](#)
- [Combining Centralized and Distributed Path Computation Paradigms | 436](#)

Abstract hops are user-defined router groups. Similar to real-hop constraints that use a sequence of individual routers, a sequence of abstract hops can be used for setting up a label-switched path (LSP). The use of abstract hops provides resiliency to sequenced path constraints. The other benefits of using abstract hops include:

Specifying a Sequence of Constraint Combinations

Currently, it is possible to specify a path that can go through links that satisfy multiple attributes. Such a path constraint is called a compound constraint combination; for example, a constraint (C_i) that includes low latency links of green color and also excludes SRLG north.

However, there is no support for specifying a path with a sequence of compound constraint combinations. For example, a sequenced constraint ($C_1, C_2, C_i, \dots C_n$) that includes low latency green links, no latency blue links, and then low latency red links.

The need for such a sequenced compound constraint combination arises when there is a requirement to establish paths through a sequence of geographical regions with a different link affinity (attributes) requirement in each region. Abstract hops meet this requirement by allowing computing nodes to map each constraint combination (C_i , for example) with the user-defined group of routers—that is, the abstract hops.

Avoiding New Network Configuration on Transit Nodes

With current path constraint specification capabilities, it is possible to include or exclude links of certain attributes along an entire path; for example, excluding SRLG west from a path. However, there is no support to either conditionally exclude or include attributes, or to apply different exclude or include attributes in different parts of the path; for example, excluding SRLG west only when traversing red links.

As a workaround, a new administrative group can be created to identify all such red links that do not have SRLG west, and configure all the relevant links appropriately with that administrative group. The drawback of this approach is that configuration changes are required throughout the network to reflect the new administrative group membership.

Instead, by using abstract hops, the configuration changes can be contained on the ingress router only. At the ingress router, the constraint combination is mapped to the abstract hop, thereby meeting the aforementioned requirement without the need for any new configuration on the transit nodes.

Combining Centralized and Distributed Path Computation Paradigms

Traffic engineering of MPLS paths can be achieved by distributed computing or with a centralized controller for computing paths. A combination of both the computation types is called the hybrid computation paradigm. The key feature of the hybrid computation approach is the ability of the centralized controller—referred to as a Path Computation Element (PCE)—to loosely specify the path computation directives, per path, to the ingress router—referred to as a Path Computation Client (PCC)—and the ability of the ingress router to use it as input for path computation.

A sequence of abstract hops serves the purpose of acting as the guideline from the centralized controller. Abstract hops provide the flexibility to the controller to weave into the path constraint and attributes. This also enables the controller to build in the element of sequence in the constraint. The controller does not have to specify each hop the path needs to take, leaving room for the ingress router to act within the limits of the guideline or directive.

[Table 11 on page 436](#) lists the key features of the hybrid computation paradigm and provides a comparison of this approach with the current path computation methods.

Table 11: Hybrid Computation for Abstract Hops

Features	Distributed Constrained Shortest Path First	Centralized Constrained Shortest Path First	Hybrid Constrained Shortest Path First
React to frequent changes in a large network	Yes		Yes
Sophisticated path computation with global view		Yes	Yes
Incorporation of business logic in path computation		Yes	Yes
Resilience (no single point of failure)	Yes		Yes
Predictability		Yes	Yes
React to network load in (close to) real time	Yes		Yes
Field tested (versus early adoption)	Yes		Yes

Junos OS Implementation of Abstract Hops

IN THIS SECTION

- [Defining Abstract Hops | 437](#)
- [Using Abstract Hops in Path Constraint | 441](#)
- [Path Computation and Backtracking | 445](#)
- [Sample Backtracking | 445](#)

The order-aware abstract hops feature is introduced in Junos OS Release 17.1. The following sections describe the implementation of abstract hops in Junos OS:

Defining Abstract Hops

An abstract hop is a group of routers that users can define to be used in setting up a label-switched path (LSP). The user can control which routers to include in the group by defining a logical combination of heterogeneous link attributes or constraints called constituent attributes. The routers with links that satisfy the defined constituent attributes make it to the group of routers representing the abstract hop.

The mapping of constituent attributes with the abstract hop is local to the computing node or the ingress of the LSP being setup. As a result, abstract hops do not have associated interior gateway protocol updates or signaling protocol extensions, and implementing abstract hops in a network does not require new configuration on the transit nodes.

A constituent list enables defining of a set of constituent traffic engineering attributes, that is identified by a user-defined name. Constituent lists are used in an abstract hop definition by using any of the following configuration statements:

- **include-any-list**—Link satisfies the constituent-list if any of the specified constituent attributes are true for the link.
- **include-all-list**—Link satisfies the constituent-list if all the specified constituent attributes are true for the link.
- **exclude-all-list**—Link satisfies the constituent-list if none of the specified constituent attributes are true for the link.
- **exclude-any-list**—Link satisfies the constituent-list if at least one of the specified constituent attributes is not true for the link.

An abstract hop is defined as a logical combination of constituent-list references that can belong to any of the aforementioned categories. To achieve this, logical operators **AND** and **OR** are included in the abstract hop definition, and applied to the constituent list.

- **OR**—At least one of the constituent-list references in the abstract hop definition must be satisfied by a link for the attached node to be part of the abstract hop.
- **AND**—All of the constituent-list references in the abstract hop definition must be satisfied by a link for the attached node to be part of the abstract hop.

Sample Abstract Hop Definition

Taking as an example, the definition of abstract hops hopA is as follows:

Abstract hops hopA must include all routers whose emanating links satisfy the logical combination of the following link attributes, respectively:

- **hopA**—((administrative group red && Srlg south) || (administrative group green || Srlg north)), where:
 - *administrative group red* and *Srlg south* belong to include-all constituent list (listA1, in this example).
 - *administrative group green* and *Srlg north* belong to include-any constituent list (listA2, in this example).
 - || is the OR operator.

The configuration for abstract hops hopA is as follows:

• **hopA configuration**

```
[edit protocols mpls]
Constituent-list listA1 {
  administrative-group red;
  Srlg south;
}
Constituent-list listA2 {
  administrative-group green;
  Srlg north;
}
Abstract-hop hopA{
  Operator OR;
  Constituent-list listA1 include-all-list;
  Constituent-list listA2 include-any-list;
}
```

Verifying Abstract Hop Configuration

The **show mpls abstract hop membership** *<abstract hop name>* command is used to view members of an abstract hop. The command output provides the abstract hop to traffic engineering database node mapping.

```
user@host> show mpls abstract-hop-membership
```

```
Abstract hop: hop1A
  Credibility: 0
Address: 128.102.165.105
Address: 128.102.166.237
Address: 128.102.168.0
Address: 128.102.173.123
```

```
Abstract hop: hopB
  Credibility: 0
Address: 128.102.160.211
Address: 128.102.165.5
Address: 128.102.166.237
Address: 128.102.172.157
Address: 128.102.172.196
```

Here, the output field **Credibility** indicates the credibility associated with interior gateway protocol in use.

The output of the **show ted database extensive local** command provides the view captured in traffic engineering database. A keyword **local** is added to indicate that the output would include any local instrumentation. The command output shows the abstract hop as an attribute of links that satisfy the associated logical combination of link attributes.

```
user@host> show ted database extensive local
```

```
TED database: 0 ISIS nodes 8 INET nodes
NodeID: 128.102.173.123
Type: Rtr, Age: 3098 secs, LinkIn: 4, LinkOut: 3
Protocol: OSPF(0.0.0.0)
To: 128.102.168.0, Local: 1.3.0.1, Remote: 1.3.0.2
Local interface index: 332, Remote interface index: 0
Color: 0x2 green
Abstract hops: hopA
Metric: 1
Static BW: 1000Mbps
Reservable BW: 1000Mbps
Available BW [priority] bps:
    [0] 970Mbps    [1] 970Mbps    [2] 970Mbps    [3] 970Mbps
    [4] 970Mbps    [5] 970Mbps    [6] 970Mbps    [7] 970Mbps
```




```
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 970Mbps      [1] 970Mbps      [2] 970Mbps      [3] 970Mbps
    [4] 970Mbps      [5] 970Mbps      [6] 970Mbps      [7] 970Mbps
To: 128.102.165.105, Local: 1.1.0.1, Remote: 1.1.0.2
Local interface index: 330, Remote interface index: 0
Srlg: south

Abstract hops: hopB
Metric: 1
Static BW: 1000Mbps
Reservable BW: 1000Mbps
Available BW [priority] bps:
  [0] 960Mbps      [1] 960Mbps      [2] 960Mbps      [3] 960Mbps
  [4] 960Mbps      [5] 960Mbps      [6] 960Mbps      [7] 960Mbps
Interface Switching Capability Descriptor(1):
  Switching type: Packet
  Encoding type: Packet
  Maximum LSP BW [priority] bps:
    [0] 960Mbps      [1] 960Mbps      [2] 960Mbps      [3] 960Mbps
    [4] 960Mbps      [5] 960Mbps      [6] 960Mbps      [7] 960Mbps
```

Abstract hop hopA is for low latency AND SRLG west, and abstract hop hopB is for excluding SRLG west. [Figure 29 on page 440](#) displays the ingress view of these abstract hops.

Figure 29: Ingress View of Abstract Hops



Using Abstract Hops in Path Constraint

The user associates a unique identifier with each abstract hop definition. This identifier is used for referring to the abstract hop in the path constraint. A sequence of abstract hops can be specified as the path constraint, similar to how real IP hops are used. The path constraint could also be a sequence of abstract hops interleaved by real IP hops.

Using abstract hops or real hops in a path constraint requires more than one Constrained Shortest Path First pass to the destination, typically one pass per hop. When real hops are provided as the path constraint, the constraint computation involves as many passes as the number of hops in the path constraint, where each pass ends on reaching a hop in the constraint list. The starting point for each pass is the destination of the previous pass, with the first pass using the ingress router as the start.

Alternatively, when path constraint uses strict or loose abstract hops, constraint computation comprises passes where each pass processes the subsequent abstract hop in the constraint list. In such a case, more than one node qualifies to be the destination for the pass. The set of nodes is called the viable router set for the pass.

An abstract hop traverses member nodes by using the following:

- Links that satisfy the logical combination of defined constituent attributes
- Any kind of links

The means of abstract hops traversing the member nodes is controlled by the use of the abstract hop qualifiers—strict, loose, and loose-link—in defining the path constraint. Taking for example, abstract hop hopA is processed differently with different qualifiers:

- **Strict**—After the last processed hop in the constraint list, the path traverses only links or nodes having membership of abstract hop hopA, before reaching a node with hopA's membership that is a feasible starting point for processing the next abstract hop.
- **Loose**—After the last processed hop in the constraint list, the path can traverse any real nodes that do not have abstract hop membership of hopA, before reaching a node with abstract hop membership hopA, which is a feasible starting point for processing the next abstract hop.
- **Loose-link**—After the last processed hop in the constraint list, the path can traverse any real nodes that do not have abstract hop membership of hopA, before reaching a node with abstract hop membership hopA, which is a feasible starting point for processing the next abstract hop. But the path should have traversed at least one link of abstract hop hopA membership in the course of the same.

In other words, the abstract hop of type loose-link is said to be processed only if any of the viable routers in the constraint is reachable through a link of associated abstract hop membership.

Sample Abstract Hops Specification

[Table 12 on page 442](#) provides sample use case for using abstract hops in path constraints.

Table 12: Using Abstract Hops in Path Constraints

Purpose of Path Constraint	Abstract Hop Qualifier	Configuration	Viable Router Set	Affinity
Traverse nodes that are members of hopA taking only links that satisfy hopA.	Strict	[edit protocols mpls] Path path_hopA_s { hopA abstract strict; }	All members of abstract hopA. That is, A1, A2...An.	hopA (pick only links that satisfy abstract hopA).
Traverse nodes that are members of hopA but not necessarily links that satisfy hopA	Loose	[edit protocols mpls] Path path_hopA_l { hopA abstract loose; }	All members of abstract hopA. That is, A1, A2...An.	None (any kind of links).
Traverse nodes that are members of hopA by taking at least one link that satisfies hopA.	Loose-link NOTE: The loose-link qualifier is viewed as loose followed by strict for the same abstract hop. In other words, hopA loose-link is the same as hopA loose and hopA strict.	[edit protocols mpls] Path path_hopA_ll { hopA abstract loose-link; }	In this case, there are two computation passes associated with hopA in the path constraint. The viable router set for both passes is: All members of abstract hopA. That is, A1, A2...An. NOTE: During path computation, a router is traversed only once.	In this case, there are two computation passes associated with hopA in the path constraint. The affinity for the two passes is: <ul style="list-style-type: none"> • Pass 1—None (any kind of links). • Pass 2—hopA (pick only links that satisfy abstract hopA).

Table 12: Using Abstract Hops in Path Constraints (*continued*)


Purpose of Path Constraint	Abstract Hop Qualifier	Configuration	Viable Router Set	Affinity
Traverse nodes that are members of hopA, taking only links that satisfy hopA, followed by nodes that are members of hopB taking only links that satisfy hopB.	Strict	[edit protocols mpls] Path path_hopA_hopB_s { hopA abstract strict; hopB abstract strict; }	<ul style="list-style-type: none"> hopA—Intersection of member set of hopA and hopB. <p>NOTE: When an abstract hop is followed by a strict abstract hop, the intersection of the two member sets is considered as viable router set.</p> <ul style="list-style-type: none"> hopB—All members of abstract hopB. That is, B1, B2...Bn. 	<ul style="list-style-type: none"> hopA—hopA (pick only links that satisfy abstract hopA). hopB—hopB (pick only links that satisfy abstract hopB).
Traverse nodes that are members of hopA taking only links that satisfy hopA, followed by nodes that are members of hopB taking any kind of links.	Strict and loose	[edit protocols mpls] Path path_hopA_s_hopB_l { hopA abstract strict; hopB abstract loose; }	<ul style="list-style-type: none"> hopA—All members of abstract hopA. That is, A1, A2...An. hopB—All members of abstract hopB. That is, B1, B2...Bn. 	<ul style="list-style-type: none"> hopA—hopA (pick only links that satisfy abstract hopA). hopB—None (pick any links).
Traverse nodes that are members of hopA by taking any kinds of links, followed by nodes that are members of hopB taking any kind of links.	Loose	[edit protocols mpls] Path path_hopA_l_hopB_l { hopA abstract loose; hopB abstract loose; }	<ul style="list-style-type: none"> hopA—All members of abstract hopA. That is, A1, A2...An. hopB—All members of abstract hopB. That is, B1, B2...Bn. 	None (pick any links).

Table 12: Using Abstract Hops in Path Constraints (continued)

Purpose of Path Constraint	Abstract Hop Qualifier	Configuration	Viable Router Set	Affinity
Traverse nodes that are members of hopA by taking any kinds of links, followed by nodes that are members of hopB taking only links that satisfy hopB.	Loose and strict	[edit protocols mpls Path path_hopA_l_hopB_s { hopA abstract loose; hopB abstract strict; }]	<ul style="list-style-type: none">• hopA—Intersection of the members of hopA and hopB. When an abstract hop is followed by a strict abstract hop, the intersection of the two member sets is considered as viable router set.• hopB—All members of abstract hopB. That is, B1, B2...Bn.	<ul style="list-style-type: none">• hopA—None (pick any links).• hopB—hopB (pick only links that satisfy abstract hopB).

Figure 30 on page 444 displays path constraints for abstract hops hopA, hopB, and hopC with loose, strict, and loose abstract hop qualifiers, respectively.

Figure 30: Sample Path Constraints for Abstract Hops



The Constrained Shortest Path First passes for the abstract hops are as follows:

- Pass 1 associated with hopA
 - Viable routers—Routers R1 and R2 (intersection of hopA and hopB, as hopB is a strict abstract hop).
 - Affinity—None (as hopA is loose).

- Pass 2 associated with hopB
 - Viable routers—Routers R1, R2, R3, and R4
 - Affinity—Pick only hopB-compliant links (as hopB is a strict abstract hop).
- Pass 3 associated with hopC
 - Viable routers—Routers R5, R6, R7, and the egress router.
 - Affinity—None (as hopC is a loose abstract hop).

Path Computation and Backtracking

In each Constrained Shortest Path First pass, when the nearest router from a viable router set is reached using links satisfying the affinity figured for the pass, the abstract hop associated with the pass is said to be processed. The viable router thus reached serves as the start for the next constraint pass. If any constraint pass fails, and it is not the one with the ingress router as start router, then the pass is backtracked to the previous pass and the process is repeated.

Sample Backtracking

When a Constrained Shortest Path First pass p (other than the first one) fails, the exit router of the previous pass ($p - 1$) that served as start for the current pass p is disqualified in the viable router set of the previous pass ($p - 1$). Then the previous pass ($p - 1$) is re-executed to find the next best exit router or destination for the pass $p - 1$ from the viable router set.

The router thus determined serves as the new start router for the pass p . This procedure is repeated as long as there are failures and there are viable routers that are not explored.

The **show mpls lsp abstract-hop-computation name *lsp-name*** command provides the various computation passes involved per LSP and the qualifying exit routers for each pass. The command output also gives the affinity per pass, and shows the current start router chosen for the pass. For each viable router, the state of backtracking is displayed, where it can be either valid or disqualified.

```
user@host> show mpls lsp abstract-computation
```

```
Path computation using abstract hops for LSP: lsp1
Path type: Primary, Path name: path1

Credibility: 0, Total no of CSPF passes: 2
CSPF pass no: 0 Start address of the pass: 128.102.173.123
Affinity: hopA
CSPF pass no: 1 Start address of the pass: 0.0.0.0
Destination: 128.102.172.157, , State: VALID

Path type: Standby, Path name: path2
```



```

Credibility: 0, Total no of CSPF passes: 3
CSPF pass no: 0 Start address of the pass: 128.102.173.123
Destination: 128.102.166.237, , State: VALID
Affinity: hopA
CSPF pass no: 1 Start address of the pass: 128.102.166.237
Destination: 128.102.160.211, , State: VALID
Destination: 128.102.165.5, , State: VALID
Destination: 128.102.166.237, , State: VALID
Destination: 128.102.172.157, , State: VALID
Destination: 128.102.172.196, , State: VALID
Affinity: hopB
CSPF pass no: 2 Start address of the pass: 128.102.172.196
Destination: 128.102.172.157, , State: VALID

```

The output field **Credibility** indicates the credibility associated with the interior gateway protocol in use.

RELATED DOCUMENTATION

[Example: Configuring Abstract Hops for MPLS LSPs | 446](#)

Example: Configuring Abstract Hops for MPLS LSPs

IN THIS SECTION

- [Requirements | 447](#)
- [Overview | 447](#)
- [Configuration | 449](#)
- [Verification | 465](#)

This example shows how to configure abstract hops for MPLS label-switched paths (LSPs). Abstract hops combine the key features of existing traffic engineering constraints that enables the user to specify an order-aware and resilient path constraint for MPLS LSPs.

Requirements

This example uses the following hardware and software components:

- Six devices that can be a combination of M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, T Series Core Routers, and PTX Series Packet Transport Routers.
- Junos OS Release 17.1 or later running on all the devices.

Before you begin:

- Configure the device interfaces.
- Configure the device router ID and assign an autonomous system (AS) number.
- Configure RSVP on all the devices.
- Configure OSPF or any other interior gateway protocol on all the devices.
- Configure administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs) on all the devices.

Overview

Junos OS Release 17.1 introduces abstract hops, which are user-defined router clusters or groups. Similar to the sequence of real-hop constraints (strict or loose), a sequence of abstract hops can be used for setting up a label-switched path (LSP). A path can use a combination of real and abstract hops as constraints.

An abstract hop is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and SRLGs, along with the ordering property of real hops. As a result, when a sequence of abstract hops is used in a path constraint, ordering is achieved among the groups of routers that meet a logical combination of link or node attributes called constituent attributes.

To configure abstract hops:

- Create constituent lists with constituent traffic engineering attributes by including the **constituent-list list-name** statement at the **[edit protocols mpls]** hierarchy level.
- Include the constituent lists in the abstract hop definition at the **[edit protocols mpls abstract-hop abstract-hop-name]** hierarchy level.
- Define path constraints that use abstract hops at the **[edit protocols mpls path path-name]** hierarchy level.

Take the following guidelines under consideration when configuring abstract hops for MPLS LSPs:

- Abstract hops are supported only in the master routing instance of a device.
- IPv6 destinations are not supported in abstract hop constraints (only IPv4 destinations work).
- Abstract hops can be strict or loose constraints.

- Abstract hops support in Junos OS Release 17.1 is provided only for intra-area MPLS LSPs and not for inter-domain, or inter-area LSPs.
- Abstract hop constraints is enabled for regular point-to-point LSPs only. Other types of MPLS LSPs, such as point-to-multipoint LSPs, externally controlled bidirectional LSPs, dynamic container LSPs, RSVP automesh LSPs, and inter-area LSPs are not supported with abstract hops configuration.
- Abstract hops do not enable computation of overall shortest path for LSPs.
- An abstract hop must not be referred to more than once in the same path constraint.
- Abstract hop constraint specifications do not affect the support for Graceful Routing Engine switchover (GRES), unified in-service software upgrade (ISSU), and nonstop routing (NSR).
- Abstract hop constraint specifications do not affect overall network performance. However, the time taken for constrained shortest path first computation increases with abstract hop configuration. The setup time for an abstract hop LSP is more than the time taken to set up an LSP without abstract hop configuration.

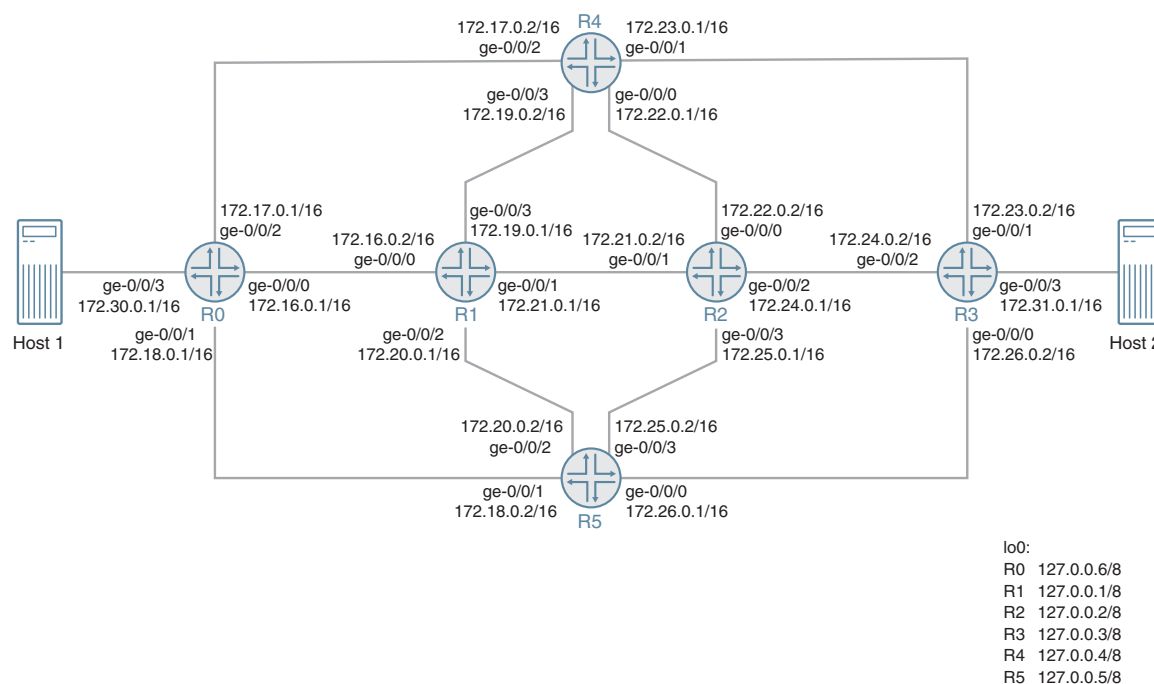
Topology

[Figure 31 on page 449](#) illustrates a sample network topology configured with abstract hops. Devices R0 and R3 are each connected to hosts (Host 1 and Host 2). Devices R4 and R5 are each connected to Devices R0, R1, R2, and R3. Devices R1 and R2 are also directly connected to each other.

Devices R0 and R3 are configured under the same autonomous system—AS 64496. An MPLS LSP is configured from Device R0 through Device R3 with one primary path and two secondary paths (standby and nonstandby secondary paths).

Four constituent lists—c1, c2, c3, and c4—are created using three SRLGs (g1, g2, and g3), three administrative groups (green, blue, and red), and one extended administrative group (gold). Three abstract hops (ah1, ah2, and ah3) are defined using the configured constituent lists, and are specified as path constraints. Abstract hop ah1 is specified as constraint for the primary path, while abstract hops ah2 and ah3 are specified as constraints for the secondary standby path and the secondary nonstandby path, respectively.

Figure 31: Configuring Abstract Hop Path Constraint



Configuration

CLI Quick Configuration

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

Device R0

```
set chassis network-services ip
set interfaces ge-0/0/0 unit 0 family inet address 172.16.0.1/16
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 172.18.0.1/16
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 172.17.0.1/16
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 172.30.0.1/16
set interfaces lo0 unit 0 family inet address 127.0.0.6/8
set routing-options srlg g1 srlg-value 100
set routing-options srlg g1 srlg-cost 1000
```



```
set routing-options srlg g2 srlg-value 200
set routing-options srlg g2 srlg-cost 2000
set routing-options srlg g3 srlg-value 300
set routing-options srlg g3 srlg-cost 3000
set routing-options administrative-groups-extended-range minimum 50000
set routing-options administrative-groups-extended-range maximum 60000
set routing-options administrative-groups-extended gold group-value 50000
set routing-options router-id 127.0.0.6
set routing-options autonomous-system 64496
set routing-options forwarding-table export test
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface ge-0/0/0.0 bandwidth 80m
set protocols rsvp interface ge-0/0/2.0 bandwidth 200m
set protocols rsvp interface ge-0/0/1.0 bandwidth 500m
set protocols mpls administrative-groups green 0
set protocols mpls administrative-groups blue 1
set protocols mpls administrative-groups red 2
set protocols mpls label-switched-path R0-R31 to 127.0.0.3
set protocols mpls label-switched-path R0-R31 primary prim
set protocols mpls label-switched-path R0-R31 secondary stdby standby
set protocols mpls label-switched-path R0-R31 secondary nonstdby
set protocols mpls path path_primary 172.16.0.2 strict
set protocols mpls path path_primary 172.21.0.2 strict
set protocols mpls path path_primary 172.24.0.2 strict
set protocols mpls path path_ter_nonstdby 172.18.0.1 strict
set protocols mpls path path_ter_nonstdby 172.26.0.2 strict
set protocols mpls path path_sec_stdby 172.17.0.2 strict
set protocols mpls path path_sec_stdby 172.23.0.2 strict
set protocols mpls path prim ah1 abstract
set protocols mpls path prim ah1 strict
set protocols mpls path stdby ah2 abstract
set protocols mpls path stdby ah2 strict
set protocols mpls path nonstdby ah3 abstract
set protocols mpls path nonstdby ah3 strict
set protocols mpls constituent-list c1 srlg g1
set protocols mpls constituent-list c1 administrative-group green
set protocols mpls constituent-list c2 administrative-group green
set protocols mpls constituent-list c2 administrative-group-extended gold
set protocols mpls constituent-list c3 srlg g2
set protocols mpls constituent-list c3 administrative-group red
set protocols mpls constituent-list c3 administrative-group-extended gold
```



```

set protocols mpls constituent-list c4 srlg g3
set protocols mpls constituent-list c4 administrative-group blue
set protocols mpls constituent-list c4 administrative-group-extended gold
set protocols mpls abstract-hop ah1 operator AND
set protocols mpls abstract-hop ah1 constituent-list c1 include-all-list
set protocols mpls abstract-hop ah1 constituent-list c2 include-all-list
set protocols mpls abstract-hop ah2 operator AND
set protocols mpls abstract-hop ah2 constituent-list c3 include-all-list
set protocols mpls abstract-hop ah3 operator AND
set protocols mpls abstract-hop ah3 constituent-list c4 include-all-list
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 srlg g1
set protocols mpls interface ge-0/0/0.0 administrative-group green
set protocols mpls interface ge-0/0/0.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/2.0 srlg g2
set protocols mpls interface ge-0/0/2.0 administrative-group red
set protocols mpls interface ge-0/0/2.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/1.0 srlg g3
set protocols mpls interface ge-0/0/1.0 administrative-group blue
set protocols mpls interface ge-0/0/1.0 administrative-group-extended gold
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set policy-options policy-statement test then load-balance per-packet

```

Device R1

```

set interfaces ge-0/0/0 unit 0 family inet address 172.16.0.2/16
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 172.21.0.1/16
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 172.20.0.1/16
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 172.19.0.1/16
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/8
set routing-options srlg g1 srlg-value 100
set routing-options srlg g1 srlg-cost 1000
set routing-options srlg g2 srlg-value 200

```



```

set routing-options srlg g2 srlg-cost 2000
set routing-options srlg g3 srlg-value 300
set routing-options srlg g3 srlg-cost 3000
set routing-options administrative-groups-extended-range minimum 50000
set routing-options administrative-groups-extended-range maximum 60000
set routing-options administrative-groups-extended gold group-value 50000
set routing-options router-id 127.0.0.1
set protocols rsvp interface fxp0.0 disable
set protocols mpls administrative-groups green 0
set protocols mpls administrative-groups blue 1
set protocols mpls administrative-groups red 2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 srlg g1
set protocols mpls interface ge-0/0/0.0 administrative-group green
set protocols mpls interface ge-0/0/0.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/1.0 srlg g1
set protocols mpls interface ge-0/0/1.0 administrative-group green
set protocols mpls interface ge-0/0/1.0 administrative-group-extended gold
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device R2

```

set interfaces ge-0/0/0 unit 0 family inet address 172.22.0.2/16
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 172.21.0.2/16
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 172.24.0.1/16
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 172.25.0.1/16
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.2/8
set routing-options srlg g1 srlg-value 100
set routing-options srlg g1 srlg-cost 1000
set routing-options srlg g2 srlg-value 200
set routing-options srlg g2 srlg-cost 2000
set routing-options srlg g3 srlg-value 300
set routing-options srlg g3 srlg-cost 3000

```



```

set routing-options administrative-groups-extended-range minimum 50000
set routing-options administrative-groups-extended-range maximum 60000
set routing-options administrative-groups-extended gold group-value 50000
set routing-options router-id 127.0.0.2
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls administrative-groups green 0
set protocols mpls administrative-groups blue 1
set protocols mpls administrative-groups red 2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/1.0 srlg g1
set protocols mpls interface ge-0/0/1.0 administrative-group green
set protocols mpls interface ge-0/0/1.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/2.0 srlg g1
set protocols mpls interface ge-0/0/2.0 administrative-group green
set protocols mpls interface ge-0/0/2.0 administrative-group-extended gold
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device R3

```

set interfaces ge-0/0/0 unit 0 family inet address 172.26.0.2/16
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 172.23.0.2/16
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 172.24.0.2/16
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 172.31.0.1/16
set interfaces lo0 unit 0 family inet address 127.0.0.3/8
set routing-options srlg g1 srlg-value 100
set routing-options srlg g1 srlg-cost 1000
set routing-options srlg g2 srlg-value 200
set routing-options srlg g2 srlg-cost 2000
set routing-options srlg g3 srlg-value 300
set routing-options srlg g3 srlg-cost 3000
set routing-options administrative-groups-extended-range minimum 50000
set routing-options administrative-groups-extended-range maximum 60000
set routing-options administrative-groups-extended gold group-value 50000

```



```

set routing-options router-id 127.0.0.3
set routing-options autonomous-system 64496
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls administrative-groups green 0
set protocols mpls administrative-groups blue 1
set protocols mpls administrative-groups red 2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/2.0 srlg g1
set protocols mpls interface ge-0/0/2.0 administrative-group green
set protocols mpls interface ge-0/0/2.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/1.0 srlg g2
set protocols mpls interface ge-0/0/1.0 administrative-group red
set protocols mpls interface ge-0/0/1.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/0.0 srlg g3
set protocols mpls interface ge-0/0/0.0 administrative-group blue
set protocols mpls interface ge-0/0/0.0 administrative-group-extended gold
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device R4

```

set interfaces ge-0/0/0 unit 0 family inet address 172.22.0.1/16
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 172.23.0.1/16
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 172.17.0.2/16
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 172.19.0.2/16
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.4/32
set routing-options srlg g1 srlg-value 100
set routing-options srlg g1 srlg-cost 1000
set routing-options srlg g2 srlg-value 200
set routing-options srlg g2 srlg-cost 2000
set routing-options srlg g3 srlg-value 300
set routing-options srlg g3 srlg-cost 3000
set routing-options administrative-groups-extended-range minimum 50000

```



```

set routing-options administrative-groups-extended-range maximum 60000
set routing-options administrative-groups-extended gold group-value 50000
set routing-options router-id 127.0.0.4
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls administrative-groups green 0
set protocols mpls administrative-groups blue 1
set protocols mpls administrative-groups red 2
set protocols mpls icmp-tunneling
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/2.0 srlg g2
set protocols mpls interface ge-0/0/2.0 administrative-group red
set protocols mpls interface ge-0/0/2.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/1.0 srlg g2
set protocols mpls interface ge-0/0/1.0 administrative-group red
set protocols mpls interface ge-0/0/1.0 administrative-group-extended gold
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Device R5

```

set interfaces ge-0/0/0 unit 0 family inet address 172.26.0.1/16
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 172.18.0.2/16
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 172.20.0.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 172.25.0.2/16
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.5/8
set routing-options srlg g1 srlg-value 100
set routing-options srlg g1 srlg-cost 1000
set routing-options srlg g2 srlg-value 200
set routing-options srlg g2 srlg-cost 2000
set routing-options srlg g3 srlg-value 300
set routing-options srlg g3 srlg-cost 3000
set routing-options administrative-groups-extended-range minimum 50000
set routing-options administrative-groups-extended-range maximum 60000

```



```

set routing-options administrative-groups-extended gold group-value 50000
set routing-options router-id 127.0.0.5
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls administrative-groups green 0
set protocols mpls administrative-groups blue 1
set protocols mpls administrative-groups red 2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/1.0 srlg g3
set protocols mpls interface ge-0/0/1.0 administrative-group blue
set protocols mpls interface ge-0/0/1.0 administrative-group-extended gold
set protocols mpls interface ge-0/0/0.0 srlg g3
set protocols mpls interface ge-0/0/0.0 administrative-group blue
set protocols mpls interface ge-0/0/0.0 administrative-group-extended gold
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Step-by-Step Procedure

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

To configure Device R0:

1. Enable enhanced IP network services on Device R0.

```

[edit chassis]
user@R0# set network-services ip

```

2. Configure the interfaces on Device R0, including the loopback interface.

```

[edit interfaces]
user@R0# set ge-0/0/0 unit 0 family inet address 172.16.0.1/16
user@R0# set ge-0/0/0 unit 0 family mpls
user@R0# set ge-0/0/1 unit 0 family inet address 172.18.0.1/16
user@R0# set ge-0/0/1 unit 0 family mpls
user@R0# set ge-0/0/2 unit 0 family inet address 172.17.0.1/16
user@R0# set ge-0/0/2 unit 0 family mpls
user@R0# set ge-0/0/3 unit 0 family inet address 172.30.0.1/16

```



```
user@R0# set lo0 unit 0 family inet address 127.0.0.6/8
```

3. Assign the router ID and autonomous system number for Device R0.

```
[edit routing-options]
user@R0# set router-id 127.0.0.6
user@R0# set autonomous-system 64496
```

4. Configure the SRLG definitions.

```
[edit routing-options]
user@R0# set srlg g1 srlg-value 100
user@R0# set srlg g1 srlg-cost 1000
user@R0# set srlg g2 srlg-value 200
user@R0# set srlg g2 srlg-cost 2000
user@R0# set srlg g3 srlg-value 300
user@R0# set srlg g3 srlg-cost 3000
```

5. Configure the extended administrative group definitions.

```
[edit routing-options]
user@R0# set administrative-groups-extended-range minimum 50000
user@R0# set administrative-groups-extended-range maximum 60000
user@R0# set administrative-groups-extended gold group-value 50000
```

6. Configure the administrative group definitions.

```
[edit protocols]
user@R0# set mpls administrative-groups green 0
user@R0# set mpls administrative-groups blue 1
user@R0# set mpls administrative-groups red 2
```

7. Configure MPLS on all the interfaces of Device R0, excluding the management interface.

```
[edit protocols]
user@R0# set mpls interface all
user@R0# set mpls interface fxp0.0 disable
```


8. Assign the interfaces of Device R0 with the configured traffic engineering attributes.

```
[edit protocols]
user@R0# set mpls interface ge-0/0/0.0 srlg g1
user@R0# set mpls interface ge-0/0/0.0 administrative-group green
user@R0# set mpls interface ge-0/0/0.0 administrative-group-extended gold
user@R0# set mpls interface ge-0/0/2.0 srlg g2
user@R0# set mpls interface ge-0/0/2.0 administrative-group red
user@R0# set mpls interface ge-0/0/2.0 administrative-group-extended gold
user@R0# set mpls interface ge-0/0/1.0 srlg g3
user@R0# set mpls interface ge-0/0/1.0 administrative-group blue
user@R0# set mpls interface ge-0/0/1.0 administrative-group-extended gold
```

9. Configure an LSP connecting Device R0 with Device R3, and assign primary and secondary path attributes to the LSP.

```
[edit protocols]
user@R0# set mpls label-switched-path R0-R31 to 127.0.0.3
user@R0# set mpls label-switched-path R0-R31 primary prim
user@R0# set mpls label-switched-path R0-R31 secondary stdby standby
user@R0# set mpls label-switched-path R0-R31 secondary nonstdby
```

10. Define the primary and secondary paths for the R0-R31 LSP.

```
[edit protocols]
user@R0# set mpls path path_primary 172.16.0.2 strict
user@R0# set mpls path path_primary 172.21.0.2 strict
user@R0# set mpls path path_primary 172.24.0.2 strict
user@R0# set mpls path path_ter_nonstdby 172.18.0.1 strict
user@R0# set mpls path path_ter_nonstdby 172.26.0.2 strict
user@R0# set mpls path path_sec_stdby 172.17.0.2 strict
user@R0# set mpls path path_sec_stdby 172.23.0.2 strict
```

11. Create constituent lists with constituent traffic engineering attributes for abstract-hop definitions.

```
[edit protocols]
user@R0# set mpls constituent-list c1 srlg g1
user@R0# set mpls constituent-list c1 administrative-group green
user@R0# set mpls constituent-list c2 administrative-group green
user@R0# set mpls constituent-list c2 administrative-group-extended gold
user@R0# set mpls constituent-list c3 srlg g2
```



```

user@R0# set mpls constituent-list c3 administrative-group red
user@R0# set mpls constituent-list c3 administrative-group-extended gold
user@R0# set mpls constituent-list c4 srlg g3
user@R0# set mpls constituent-list c4 administrative-group blue
user@R0# set mpls constituent-list c4 administrative-group-extended gold

```

12. Define abstract hops by assigning the configured constituent lists and respective operators.

```

[edit protocols]
user@R0# set mpls abstract-hop ah1 operator AND
user@R0# set mpls abstract-hop ah1 constituent-list c1 include-all-list
user@R0# set mpls abstract-hop ah1 constituent-list c2 include-all-list
user@R0# set mpls abstract-hop ah2 operator AND
user@R0# set mpls abstract-hop ah2 constituent-list c3 include-all-list
user@R0# set mpls abstract-hop ah3 operator AND
user@R0# set mpls abstract-hop ah3 constituent-list c4 include-all-list

```

13. Define constraints for the configured paths by including abstract hop definitions.

```

[edit protocols]
user@R0# set mpls path prim ah1 abstract
user@R0# set mpls path prim ah1 strict
user@R0# set mpls path stdby ah2 abstract
user@R0# set mpls path stdby ah2 strict
user@R0# set mpls path nonstdby ah3 abstract
user@R0# set mpls path nonstdby ah3 strict

```

14. Configure RSVP on Device R0. Enable RSVP on all the interfaces of Device R0, excluding the management interface and interface connecting to Host1, and assign bandwidth values.

```

[edit protocols]
user@R0# set rsvp interface all aggregate
user@R0# set rsvp interface fxp0.0 disable
user@R0# set rsvp interface ge-0/0/0.0 bandwidth 80m
user@R0# set rsvp interface ge-0/0/2.0 bandwidth 200m
user@R0# set rsvp interface ge-0/0/1.0 bandwidth 500m

```

15. Configure OSPF on all the interfaces of Device R0, excluding the management interface, and assign traffic engineering capabilities.


```
[edit protocols]
user@R0# set ospf traffic-engineering
user@R0# set ospf area 0.0.0.0 interface all
user@R0# set ospf area 0.0.0.0 interface fxp0.0 disable
```

16. Configure a policy on Device R0 to enable load balancing on a per-packet basis.

```
[edit policy-options]
user@R0# set forwarding-table export test
```

17. Export the load-balancing policy to the forwarding table.

```
[edit policy-options]
user@R0# set policy-statement test then load-balance per-packet
```

Results

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

```
user@R0# show chassis
network-services ip;
```

```
user@R0# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 172.16.0.1/16;
    }
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 172.18.0.1/16;
    }
    family mpls;
  }
}
```



```

ge-0/0/2 {
  unit 0 {
    family inet {
      address 172.17.0.1/16;
    }
    family mpls;
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 172.30.0.1/16;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.6/8;
    }
  }
}

```

```

user@R0# show routing-options
srlg {
  g1 {
    srlg-value 100;
    srlg-cost 1000;
  }
  g2 {
    srlg-value 200;
    srlg-cost 2000;
  }
  g3 {
    srlg-value 300;
    srlg-cost 3000;
  }
}
administrative-groups-extended-range {
  minimum 50000;
  maximum 60000;
}
administrative-groups-extended {
  gold group-value 50000;
}

```



```
}
```

```
user@R0# show protocols
```

```

rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface ge-0/0/0.0 {
    bandwidth 80m;
  }
  interface ge-0/0/2.0 {
    bandwidth 200m;
  }
  interface ge-0/0/1.0 {
    bandwidth 500m;
  }
}
mpls {
  administrative-groups {
    green 0;
    blue 1;
    red 2;
  }
  label-switched-path R0-R31 {
    to 127.0.0.3;
    adaptive;
    auto-bandwidth {
      adjust-interval 300;
      adjust-threshold 5;
      minimum-bandwidth 10m;
      maximum-bandwidth 1g;
    }
    primary prim;
    secondary stdby {
      standby;
    }
    secondary nonstdby;
  }
  path path_primary {
    172.16.0.2 strict;
    172.21.0.2 strict;
    172.24.0.2 strict;
  }
}

```



```

path path_ter_nonstdby {
    172.18.0.1 strict;
    172.26.0.2 strict;
}
path path_sec_stdby {
    172.17.0.2 strict;
    172.23.0.2 strict;
}
path prim {
    ah1 abstract strict;
}
path stdby {
    ah2 abstract strict;
}
path nonstdby {
    ah3 abstract strict;
}
constituent-list c1 {
    srlg g1;
    administrative-group green;
}
constituent-list c2 {
    administrative-group green;
    administrative-group-extended gold;
}
constituent-list c3 {
    srlg g2;
    administrative-group red;
    administrative-group-extended gold;
}
constituent-list c4 {
    srlg g3;
    administrative-group blue;
    administrative-group-extended gold;
}
abstract-hop ah1 {
    operator AND;
    constituent-list {
        c1 include-all-list;
        c2 include-all-list;
    }
}
abstract-hop ah2 {
    operator AND;

```



```

    constituent-list {
        c3 include-all-list;
    }
}
abstract-hop ah3 {
    operator AND;
    constituent-list {
        c4 include-all-list;
    }
}
interface all;
interface fxp0.0 {
    disable;
}
interface ge-0/0/0.0 {
    srlg g1;
    administrative-group green;
    administrative-group-extended gold;
}
interface ge-0/0/2.0 {
    srlg g2;
    administrative-group red;
    administrative-group-extended gold;
}
interface ge-0/0/1.0 {
    srlg g3;
    administrative-group blue;
    administrative-group-extended gold;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
}

```

```

user@R0# show policy-options
policy-statement test {
    then {
        load-balance per-packet;
    }
}

```



```

    }
}

```

Verification

IN THIS SECTION

- [Verifying Abstract Hop Configuration | 465](#)
- [Verifying Abstract Hop Path Computation | 466](#)

Confirm that the configuration is working properly.

Verifying Abstract Hop Configuration

Purpose

Verify the members of the abstract hop definition on Device R0 by issuing the **show mpls abstract-hop-membership** command, which displays the abstract hop membership tables.

Action

From operational mode, run the **show mpls abstract-hop-membership** command.

```
user@R0> show mpls abstract-hop-membership
```

Abstract hop: ah1

```

    Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.1
Address: 127.0.0.2
Address: 127.0.0.3

```

Abstract hop: ah2

```

    Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.3
Address: 127.0.0.4

```

Abstract hop: ah3

```

    Credibility: 0

```



```
Address: 127.0.0.6
Address: 127.0.0.3
Address: 127.0.0.5
```

Meaning

The **show mpls abstract-hop-membership** command output provides the abstract hop to traffic engineering database node mapping. The **Credibility** field displays the credibility value associated with the interior gateway protocol in use (OSPF).

Verifying Abstract Hop Path Computation

Purpose

Verify the abstract computation preprocessing for LSPs on Device R0 by issuing the **show mpls lsp abstract-computation** command.

Action

From operational mode, run the **show mpls lsp abstract-computation** command.

```
user@R0> show mpls lsp abstract-computation
```

```
Path computation using abstract hops for LSP: R0-R31
```

```
Path type: Primary, Path name: prim
```

```
Credibility: 0, Total no of CSPF passes: 2
```

```
CSPF pass no: 0
```

```
Start address of the pass: 127.0.0.6
```

```
Destination: 127.0.0.1, State: VALID
```

```
Destination: 127.0.0.2, State: VALID
```

```
Destination: 127.0.0.3, State: VALID
```

```
Affinity: ah1
```

```
CSPF pass no: 1
```

```
Start address of the pass: 127.0.0.1
```

```
Destination: 127.0.0.3, State: VALID
```

```
Path type: Secondary, Path name: nonstdby
```

```
Path type: Standby, Path name: stdby
```

```
Credibility: 0, Total no of CSPF passes: 2
```

```
CSPF pass no: 0
```

```
Start address of the pass: 127.0.0.6
```

```
Destination: 127.0.0.3, State: VALID
```

```
Destination: 127.0.0.4, State: VALID
```

```
Affinity: ah2
```



```
CSPF pass no: 1
Start address of the pass: 127.0.0.4
Destination: 127.0.0.3, State: VALID
```

Meaning

The **show mpls lsp abstract-hop-computation** command output provides the various computation passes involved per LSP, and the qualifying exit devices for each pass. The command output also gives the affinity per pass, and shows the current start device chosen for the pass. For each viable router (device), the state of backtracking is displayed, where it can either be valid or disqualified.

The **Credibility** field indicates the credibility value associated with the interior gateway protocol in use (OSPF).

RELATED DOCUMENTATION

[Abstract Hops for MPLS LSPs Overview | 433](#)

[constituent-list | 2578](#)

[abstract-hop | 2534](#)

[show mpls abstract-hop-membership | 3170](#)

[show mpls lsp abstract-computation | 3240](#)

Configuring the Maximum Number of MPLS Labels

For interfaces that you configure for MPLS applications, you can set the maximum number of labels upon which MPLS can operate.

By default, the maximum number of labels is three. You can change the maximum to four labels or five labels for applications that require four or five labels.

Starting in Junos OS Release 19.1R1, the maximum number of labels that can be pushed by the egress Packet Forwarding Engine (PFE) can be leveraged, wherein the number of labels that can be pushed for an MPLS next hop is the number of labels the device is capable of pushing, or the [maximum-labels](#) configured under **family mpls** of the outgoing interface, whichever is smaller. This support is enabled on MX Series routers with MPC and MIC interfaces, and PTX Series routers with third-generation FPCs.

The increased label push capability is useful for features, such as segment routing traffic-engineering LSPs and RSVP-TE pop-and-forward LSPs. All existing functionality of applications using MPLS next hops continue to work with the increased label push capability. This includes:

- All OAM utilities, such as lsping, traceroute, and BFD for MPLS LSPs.
- Monitoring utilities, such as lspmon, and LM DM for MPLS LSPs.

The **show route table** and **show route forwarding-table** command outputs are enhanced to display up to 16 labels per next hop component.

For example:

```
user@host> show route table inet.3

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

11.0.0.17/32          *[SPRING-TE/8] 00:02:16, metric 1
                    > to 192.1.2.2 via ge-0/0/2.0, Push 1000115, Push 1000114,
Push 1000113, Push 1000112, Push 1000111, Push 1000110, Push 1000109, Push 1000108,
Push 1000107, Push 1000106, Push 1000105, Push 1000104, Push 1000103, Push 1000102,
Push 1000101(top)
                    to 192.1.3.2 via ge-0/0/4.0, Push 1000115, Push 1000114,
Push 1000113, Push 1000112, Push 1000111, Push 1000110, Push 1000109, Push 1000108,
Push 1000107, Push 1000106, Push 1000105, Push 1000104, Push 1000103, Push 1000102,
Push 1000101(top)
```

NOTE: When the maximum number of MPLS labels of an interface is modified, the MPLS interface is bounced. All LDP and RSVP sessions on that interface are restarted, resulting in all LSPs over that interface to flap.

For example, suppose you configure a two-tier carrier-of-carriers VPN service for customers who provide VPN service. A carrier-of-carrier VPN is a two-tiered relationship between a provider carrier (Tier 1 ISP) and a customer carrier (Tier 2 ISP). In a carrier-of-carrier VPN, the provider carrier provides a VPN backbone network for the customer carrier. The customer carrier in turn provides Layer 3 VPN service to its end customers. The customer carrier sends labeled traffic to the provider carrier to deliver it to the next hop on the other side of the provider carrier's network. This scenario requires a three-label stack: one label for the provider carrier VPN, another label for the customer carrier VPN, and a third label for the transport route.

If you add fast reroute service, the PE routers in the provider carrier's network must be configured to support a fourth label (the reroute label). If the customer carrier is using LDP as its signaling protocol and the provider carrier is using RSVP, the provider carrier must support LDP over RSVP tunnel service. This additional service requires an additional label, for a total of five labels.

To the customer carrier, the router it uses to connect to the provider carrier's VPN is a PE router. However, the provider carrier views this device as a CE router.

[Table 13 on page 469](#) summarizes the label requirements.

Table 13: Sample Scenarios for Using 3, 4, or 5 MPLS Labels

Number of Labels Required	Scenarios
3	Carrier-of-carriers VPN or a VPN with two labels and fast reroute
4	Combination of carrier-of-carriers and fast reroute
5	Carrier-of-carriers with fast reroute and the customer carrier running LDP, with the provider carrier running RSVP

To configure and monitor the maximum number of labels:

1. Specify the maximum on the logical interface. Apply this configuration to the carrier's PE routers.

```
[edit interfaces ge-0/1/3 unit 0 family mpls]
user@switch# set maximum-labels maximum-limit
```

2. Verify the configuration.

```
[edit system]
user@switch# show interfaces ge-0/1/3.0

Logical interface ge-0/1/3.0 (Index 77) (SNMP ifIndex 507)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol mpls, MTU: 1480, Maximum labels: 8
  Flags: Is-Primary
```

The command output includes the **Maximum labels: 5** field under the logical interface unit 0.

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1, the maximum number of labels that can be pushed by the egress Packet Forwarding Engine (PFE) can be leveraged, wherein the number of labels that can be pushed for an MPLS next hop is the number of labels the device is capable of pushing, or the maximum-labels configured under family mpls of the outgoing interface, whichever is smaller. This support is enabled on MX Series routers with MPC and MIC interfaces, and PTX Series routers with third-generation FPCs.

RELATED DOCUMENTATION

[maximum-labels](#) | [2679](#)

[Fast Reroute Overview](#) | [477](#)

[Tunneling LDP LSPs in RSVP LSPs Overview](#) | [928](#)

Junos VPNs Configuration Guide for a carrier-of-carriers configuration example

Configuring MPLS to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of a label-switched path (LSP). The default advertised label is label 3 (Implicit Null Label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. By enabling ultimate-hop popping, label 0 (IPv4 Explicit Null Label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

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

To configure MPLS to pop the label on the ultimate-hop router, include the **explicit-null** statement:

```
explicit-null;
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols mpls]`

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

RELATED DOCUMENTATION

[MPLS Label Overview](#) | 419

[MPLS Label Allocation](#) | 420

Advertising Explicit Null Labels to BGP Peers

For the IPv4 (**inet**) family only, BGP peers in a routing group can send an explicit NULL label for a set of connected routes (direct and loopback routes) for the inet labeled-unicast and inet6 labeled-unicast NLRI. By default, peers advertise label 3 (implicit NULL). If the **explicit-null** statement is enabled, peers advertise label 0 (explicit NULL). The explicit NULL labels ensures that labels are always present on packets traversing an MPLS network. If the implicit NULL label is used, the penultimate hop router removes the label and sends the packet as a plain IP packet to the egress router. This might cause issues in queuing the packet properly on the penultimate hop router if the penultimate hop is another vendor's router. Some other vendors queue packets based on the CoS bits in the outgoing label rather than the incoming label.

To advertise an explicit null label, include the following statements in the configuration:

```
family inet {
  labeled-unicast {
    aggregate-label {
      community community-name;
    }
    explicit-null {
      connected-only;
    }
  }
}
```

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

The **connected-only** statement is required to advertise explicit null labels.

To verify that the explicit NULL label is being advertised for connected routes, use the **show route advertising-protocol bgp *neighbor-address*** command.

RELATED DOCUMENTATION

[Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router | 1117](#)

[Configuring RSVP to Pop the Label on the Ultimate-Hop Router | 858](#)

MPLS LSP Routes

IN THIS CHAPTER

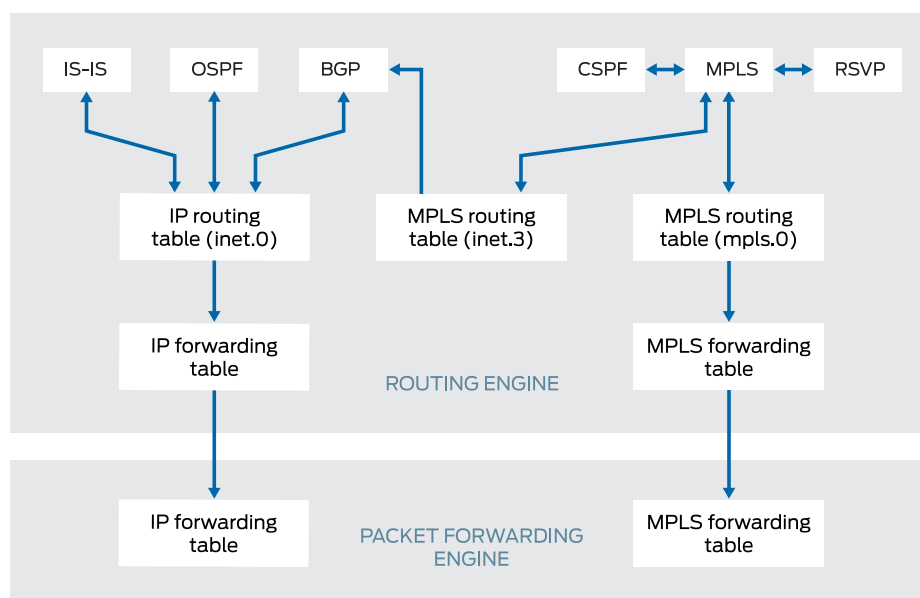
- [MPLS and Routing Tables | 474](#)
- [MPLS and Traffic Protection | 476](#)
- [Fast Reroute Overview | 477](#)
- [Configuring Fast Reroute | 479](#)
- [Detour Merging Process | 481](#)
- [Detour Computations | 482](#)
- [Fast Reroute Path Optimization | 482](#)
- [Configuring the Optimization Interval for Fast Reroute Paths | 483](#)
- [Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table | 483](#)
- [Constrained-Path LSP Computation | 485](#)
- [How CSPF Selects a Path | 486](#)
- [CSPF Path Selection Tie-Breaking | 487](#)
- [Computing CSPF Paths Offline | 488](#)
- [Configuring CSPF Tie Breaking | 489](#)
- [Disabling Constrained-Path LSP Computation | 489](#)
- [Configuring Load Balancing Based on MPLS Labels | 490](#)
- [Configuring Load Balancing Based on MPLS Labels on ACX Series Routers | 496](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes All Forwarding Decisions | 500](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and Considers Hop Constraints | 501](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and the Secondary Path Is Explicit | 501](#)
- [Path Computation for LSPs on an Overloaded Router | 503](#)
- [Computing Backup Paths for LSPs Using Fate Sharing | 504](#)
- [Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts | 504](#)
- [Enabling IGP Shortcuts | 506](#)
- [LSPs Qualified in IGP Shortcut Computations | 506](#)
- [IGP Shortcut Applications | 506](#)
- [IGP Shortcuts and Routing Tables | 507](#)

- IGP Shortcuts and VPNs | 508
- Advertising LSPs into IGPs | 508
- Selecting a Forwarding LSP Next Hop | 509
- Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes | 510
- ECMP Flow-Based Forwarding on ACX Series Routers | 511

MPLS and Routing Tables

The IGPs and BGP store their routing information in the inet.0 routing table, the main IP routing table. If the **traffic-engineering bgp** command is configured, thereby allowing only BGP to use MPLS paths for forwarding traffic, MPLS path information is stored in a separate routing table, inet.3. Only BGP accesses the inet.3 routing table. BGP uses both inet.0 and inet.3 to resolve next-hop addresses. If the **traffic-engineering bgp-igp** command is configured, thereby allowing the IGPs to use MPLS paths for forwarding traffic, MPLS path information is stored in the inet.0 routing table. (Figure 32 on page 474 and Figure 33 on page 475 illustrate the routing tables in the two traffic engineering configurations.)

Figure 32: Routing and Forwarding Tables, traffic-engineering bgp



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

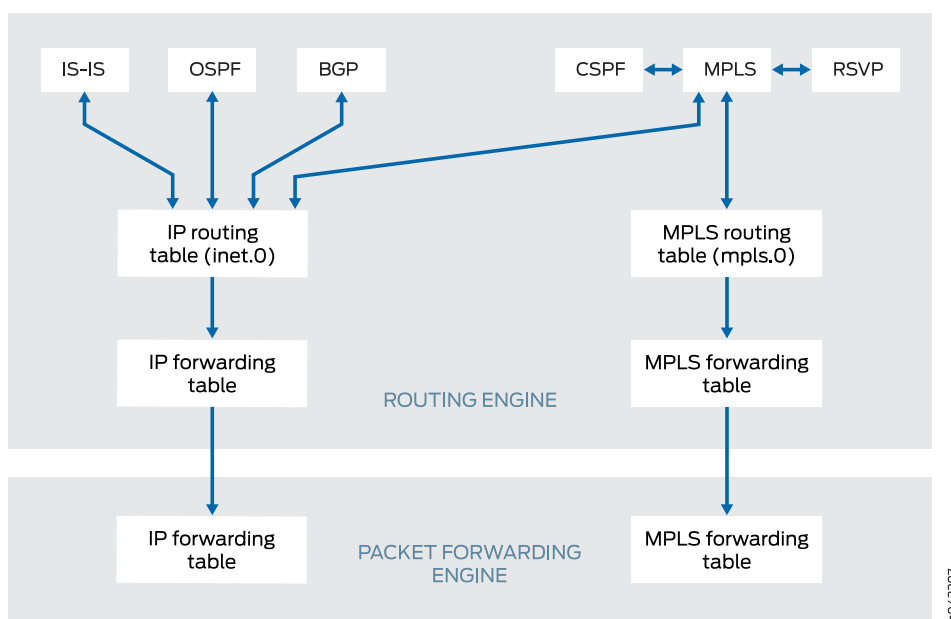
MPLS also maintains an MPLS path routing table (mpls.0), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

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

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

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

Figure 33: Routing and Forwarding Tables, traffic-engineering bgp-igp



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

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

MPLS and Traffic Protection

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

- **Standby secondary paths**—You can configure primary and secondary paths. You configure secondary paths with the **standby** statement. To activate traffic protection, you need to configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see [“Configuring Hot Standby of Secondary Paths for LSPs” on page 568](#).
- **Fast reroute**—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see [“Fast Reroute Overview” on page 477](#). For information about configuring fast reroute, see [“Configuring Fast Reroute” on page 479](#).
- **Link protection**—You can configure link protection to help ensure that traffic traversing a specific interface from one router to another can continue to reach its destination in the event that this interface fails. When link protection is configured for an interface and configured for an LSP that traverses this interface, a bypass LSP is created that handles this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination. For information about configuring link protection, see [“Configuring Link Protection on Interfaces Used by LSPs” on page 871](#).

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

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

RELATED DOCUMENTATION

[Configuring Hot Standby of Secondary Paths for LSPs | 568](#)

[Fast Reroute Overview | 477](#)

[Configuring Fast Reroute | 479](#)


[Configuring Link Protection on Interfaces Used by LSPs | 871](#)

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 34 on page 477](#) 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 is a failure in a scaled fast reroute scenario, the devices lose reachability to all the peers that were connected through the failed link. This leads to traffic interruption, as the BGP session among the devices goes down. 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 34: 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 35 on page 478](#) illustrates the detour taken when the link between Router B and Router C fails.

Figure 35: 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 34 on page 477](#) 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 36 on page 479](#), 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 36: Detours Merging into Other Detours



RELATED DOCUMENTATION

[fast-reroute](#) | 2619

[Configuring Fast Reroute](#) | 479

[MPLS Feature Support on QFX Series and EX4600 Switches](#) | 26

[Understanding Interprovider and Carrier-of-Carriers VPNs](#) | 1647

Configuring Fast Reroute

Fast reroute provides a mechanism for automatically rerouting traffic on an LSP if a node or link in an LSP fails, thus reducing the loss of packets traveling over the LSP.

To configure fast reroute on an LSP, include the **fast-reroute** statement on the ingress router (or switch):

```
fast-reroute {
  (bandwidth bps | bandwidth-percent percentage);
  (exclude [ group-names ] | no-exclude );
  hop-limit number;
  (include-all [ group-names ] | no-include-all);
  (include-any [ group-names ] | no-include-any);
}
```


You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

You do not need to configure fast reroute on the LSP's transit and egress routers (or switches). Once fast reroute is enabled, the ingress router (or switch) signals all the downstream routers (or switches) that fast reroute is enabled on the LSP, and each downstream router does its best to set up detours for the LSP. If a downstream router does not support fast reroute, it ignores the request to set up detours and continues to support the LSP. A router that does not support fast reroute will cause some of the detours to fail, but otherwise has no impact on the LSP.

NOTE: To enable PFE fast reroute, configure a routing policy statement with the **load-balance per-packet** statement at the [edit policy-options policy-statement *policy-name* then] hierarchy level on each of the routers where traffic might be rerouted. See also [“Configuring Load Balancing Across RSVP LSPs” on page 848](#).

By default, no bandwidth is reserved for the rerouted path. To allocate bandwidth for the rerouted path, include either the **bandwidth** statement or the **bandwidth-percent** statement. You can only include one of these statements at a time. If you do not include either the **bandwidth** statement or the **bandwidth-percent** statement, the default setting is to not reserve bandwidth for the detour path.

When you include the **bandwidth** statement, you can specify the specific amount of bandwidth (in bits per second [bps]) you want to reserve for the detour path. The bandwidth does not need to be identical to that allocated for the LSP.

When you specify a bandwidth percent using the **bandwidth-percent** statement, the detour path bandwidth is computed by multiplying the bandwidth percentage by the bandwidth configured for the main traffic-engineered LSP. For information about how to configure the bandwidth for a traffic-engineered LSP, see [“Configuring Traffic-Engineered LSPs” on page 1213](#).

Hop-limit constraints define how many more routers a detour is allowed to traverse compared with the LSP itself. By default, the hop limit is set to 6. For example, if an LSP traverses 4 routers, any detour for the LSP can be up to 10 (that is, 4 + 6) router hops, including the ingress and egress routers.

By default, a detour inherits the same administrative (coloring) group constraints as its parent LSP when CSPF is determining the alternate path. Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. If you specify the **include-any** statement when configuring the parent LSP, all links traversed by the alternate session must have at least one color found in the list of groups. If you specify the **include-all** statement when configuring the parent LSP, all links traversed by the alternate session must have all of the colors found in the list of groups. If you specify the **exclude** statement when configuring the parent LSP, none of the links must have a color found in the list of groups. For more

information about administrative group constraints, see [“Configuring Administrative Groups for LSPs” on page 532](#).

RELATED DOCUMENTATION

[Fast Reroute Overview | 477](#)

[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)

Detour Merging Process

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

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

- When all the path messages do not include a fast reroute or a detour object, or when the router is the egress of the LSP, no merging is required. The messages are processed according to RSVP traffic engineering.
- Otherwise, the router *must* record the path state in addition to the incoming interface. If the path messages do not share the same outgoing interface and next-hop router, the router considers them to be independent LSPs and does not merge them.
- For all the path messages that share the same outgoing interface and next-hop router, the router uses the following process to select the final LSP:
 - If only one LSP originates from this node, select it as the final LSP.
 - If only one LSP contains a fast reroute object, select it as the final LSP.
 - If there are several LSPs and some of them have a detour object, eliminate those containing a detour object from the final LSP selection process.
 - If several final LSP candidates remain (that is, there are still both detour and protected LSPs), select the LSPs with fast reroute objects.
 - If none of the LSPs have fast reroute objects, select the ones without detour objects. If all the LSPs have detour objects, select them all.

- Of the remaining LSP candidates, eliminate from consideration those that traverse nodes that other LSPs avoid.
- If several candidate LSPs still remain, select the one with the shortest explicit route object (ERO) path length. If more than one LSP has the same path length, select one randomly.
- Once the final LSP has been identified, the router must transmit only the path messages that correspond to this LSP. All other LSPs are considered merged at this node.

Detour Computations

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

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

Fast Reroute Path Optimization

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

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

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

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

Configuring the Optimization Interval for Fast Reroute Paths

You can enable path optimization for fast reroute by configuring the fast reroute optimize timer. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.

To enable fast reroute path optimization, specify the number of seconds using the `optimize-timer` option for the **fast-reroute** statement:

```
fast-reroute seconds;
```

You can include this statement at the following hierarchy levels:

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

Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table

By default, a host route toward the egress router is installed in the inet.3 or inet6.3 routing table. (The host route address is the one you configure in the **to** statement.) Installing the host route allows BGP to perform next-hop resolution. It also prevents the host route from interfering with prefixes learned from dynamic routing protocols and stored in the inet.0 or inet6.0 routing table.

Unlike the routes in the inet.0 or inet6.0 table, routes in the inet.3 or inet6.3 table are not copied to the Packet Forwarding Engine, and hence they cause no changes in the system forwarding table directly. You cannot use the **ping** or **tracert** command through these routes. The only use for inet.3 or inet6.3 is to permit BGP to perform next-hop resolution. To examine the inet.3 or inet6.3 table, use the **show route table inet.3** or **show route table inet6.3** command.

To inject additional routes into the inet.3 or inet6.3 routing table, include the **install** statement:

```
install {
  destination-prefix <active>;
}
```


You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]
- [edit protocols mpls [static-label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [static-label-switched-path](#) *lsp-name*]

The specified routes are installed as aliases into the routing table when the LSP is established. Installing additional routes allows BGP to resolve next hops within the specified prefix and to direct additional traffic for these next hops to a particular LSP.

Including the **active** option with the **install** statement installs the specified prefix into the inet.0 or inet6.0 routing table, which is the primary forwarding table. The result is a route that is installed in the forwarding table any time the LSP is established, which means you can ping or trace the route. Use this option with care, because this type of prefix is very similar to a static route.

You use alias routes for routers that have multiple addresses being used as BGP next hops, or for routers that are not MPLS capable. In either of these cases, the LSP can be configured to another MPLS capable system within the local domain, which then acts as a “border” router. The LSP then terminates on the border router and, from that router, Layer 3 forwarding takes the packet to the true next-hop router.

In the case of an interconnect, the domain’s border router can act as the proxy router and can advertise the prefix for the interconnect if the border router is not setting the BGP next hop to itself.

In the case of a point of presence (POP) that has routers that do not support MPLS, one router (for example, a core router) that supports MPLS can act as a proxy for the entire POP and can inject a set of prefixes that cover the POP. Thus, all routers within the POP can advertise themselves as interior BGP (IBGP) next hops, and traffic can follow the LSP to reach the core router. This means that normal IGP routing would prevail within the POP.

You cannot use the **ping** or **traceroute** commands on routes in the inet.3 or inet6.3 routing table.

For BGP next-hop resolution, it makes no difference whether a route is in inet.0/inet6.0 or inet.3/inet6.3; the route with the best match (longest mask) is chosen. Among multiple best-match routes, the one with the highest preference value is chosen.

NOTE: The **install *destination-prefix* active** statement is not supported on static LSPs. When the **install *destination-prefix* active** statement is configured for a static LSP, the MPLS routes do not get installed into the inet.0 routing table.

Constrained-Path LSP Computation

The Constrained Shortest Path First (CSPF) algorithm is an advanced form of the shortest-path-first (SPF) algorithm used in OSPF and IS-IS route computations. CSPF is used in computing paths for LSPs that are subject to multiple constraints. When computing paths for LSPs, CSPF considers not only the topology of the network, but also the attributes of the LSP and the links, and it attempts to minimize congestion by intelligently balancing the network load.

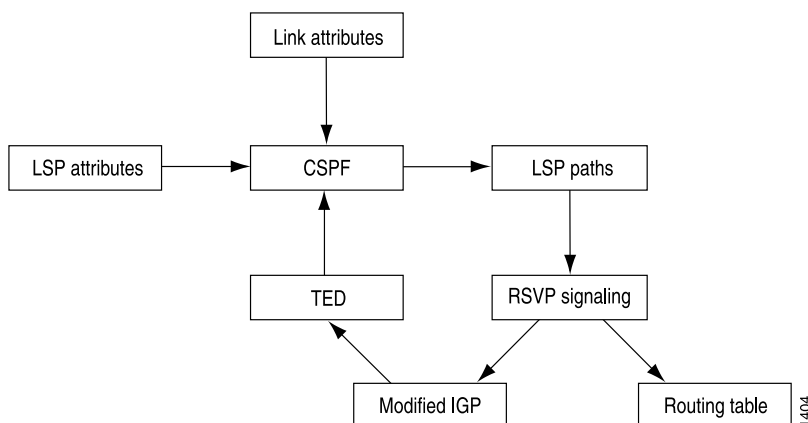
The constraints that CSPF considers include:

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

The data that CSPF considers comes from the following sources:

- Traffic engineering database—Provides CSPF with up-to-date topology information, the current reservable bandwidth of links, and the link colors. For the CSPF algorithm to perform its computations, a link-state IGP (such as OSPF or IS-IS) with special extensions is needed. For CSPF to be effective, the link-state IGP on all routers must support the special extensions. While building the topology database, the extended IGP must take into consideration the current LSPs and must flood the route information everywhere. Because changes in the reserved link bandwidth and link color cause database updates, an extended IGP tends to flood more frequently than a normal IGP. See [Figure 37 on page 486](#) for a diagram of the relationships between these components.
- Currently active LSPs—Includes all the LSPs that should originate from the router and their current operational status (up, down, or timeout).

Figure 37: CSPF Computation Process



This section discusses the following topics:

- [How CSPF Selects a Path on page 486](#)
- [CSPF Path Selection Tie-Breaking on page 487](#)
- [Computing CSPF Paths Offline on page 488](#)

How CSPF Selects a Path

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

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

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

RELATED DOCUMENTATION

[Configuring CSPF Tie Breaking | 489](#)

[CSPF Path Selection Tie-Breaking | 487](#)

CSPF Path Selection Tie-Breaking

If more than one path is still available after the CSPF rules ([“How CSPF Selects a Path” on page 486](#)) have been applied, a tie-breaking rule is applied to choose the path for the LSP. The rule used depends on the configuration. There are three tie-breaking rules:

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

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

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

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

RELATED DOCUMENTATION

[How CSPF Selects a Path | 486](#)

[Configuring CSPF Tie Breaking | 489](#)

[Configuring the Bandwidth Value for LSPs | 546](#)

[Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth | 550](#)

Computing CSPF Paths Offline

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

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

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

Configuring CSPF Tie Breaking

When selecting a path for an LSP, CSPF uses a tie-breaking process if there are several equal-cost paths. For information about how CSPF selects a path, see [“How CSPF Selects a Path” on page 486](#).

You can configure one of the following statements (you can only configure one of these statements at a time) to alter the behavior of CSPF tie-breaking:

- By default, a random tie-breaking rule for CSPF is used to select a path from the set of equal-cost paths. However, you can also explicitly configure this behavior using the **random** statement:

```
random;
```

- To prefer the path with the least-utilized links, include the **least-fill** statement:

```
least-fill;
```

- To prefer the path with the most-utilized links, include the **most-fill** statement:

```
most-fill;
```

You can include each of these statements at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name*]**
- **[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]**

RELATED DOCUMENTATION

[How CSPF Selects a Path](#) | 486

Disabling Constrained-Path LSP Computation

If the IGP is a link-state protocol (such as IS-IS or OSPF) and supports extensions that allow the current bandwidth reservation on each router's link to be reported, constrained-path LSPs are computed by default.

The Junos implementations of IS-IS and OSPF include the extensions that support constrained-path LSP computation.

- IS-IS—These extensions are enabled by default. To disable this support, include the **disable** statement at the **[edit protocols isis traffic-engineering]** hierarchy level, as discussed in the *Junos OS Routing Protocols Library*.
- OSPF—These extensions are disabled by default. To enable this support, include the **traffic-engineering** statement in the configurations of all routers running OSPF, as described in the *Junos OS Routing Protocols Library*.

If IS-IS is enabled on a router or you enable OSPF traffic engineering extensions, MPLS performs the constrained-path LSP computation by default. For information about how constrained-path LSP computation works, see [“Constrained-Path LSP Computation” on page 485](#).

Constrained-path LSPs have a greater chance of being established quickly and successfully for the following reasons:

- The LSP computation takes into account the current bandwidth reservation.
- Constrained-path LSPs reroute themselves away from node failures and congestion.

When constrained-path LSP computation is enabled, you can configure the LSP so that it is periodically reoptimized, as described in [“Optimizing Signaled LSPs” on page 541](#).

When an LSP is being established or when an existing LSP fails, the constrained-path LSP computation is repeated periodically at the interval specified by the retry timer until the LSP is set up successfully. Once the LSP is set up, no recomputation is done. For more information about the retry timer, see [“Configuring the Connection Between Ingress and Egress Routers” on page 522](#).

By default, constrained-path LSP computation is enabled. You might want to disable constrained-path LSP computation when all nodes do not support the necessary traffic engineering extensions. To disable constrained-path LSP computation, include the **no-cspf** statement:

```
no-cspf;
```

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

If you disable constrained-path LSP computation on LSPs by configuring the **no-cspf** statement and then attempt to advertise other LSPs with lower metrics than the IGPs from this router in either IS-IS or OSPF, new LSPs cannot be established.

Configuring Load Balancing Based on MPLS Labels

Load balancing occurs on a per-packet basis for MPLS flows on supported platforms. Entropy, or random distribution, is essential for the uniform distribution of packets to their next hops. By default, when load balancing is used to help distribute traffic, Junos OS employs a hash algorithm to select a next-hop address

to install into the forwarding table. Whenever the set of next hops for a destination changes, the next-hop address is reselected by means of the hash algorithm. You can configure how the hash algorithm is used to load-balance traffic across a set of equal-cost label switched paths (LSPs).

To ensure entropy for VPLS & VPWS traffic, Junos OS can create a hash based on data from the IP header and as many as three MPLS labels (the so-called top labels).

In some cases, as the number of network feature that use labels grows (such as MPLS Fast Reroute, and RFC 3107, RSVP and VPN) data in the top three labels can become static and thus not a sufficient source for entropy. Load balancing can become skewed as a result, or the incidence of out-of-order packet delivery may rise. For these cases, labels from the bottom of the label stack can be used (see Table 1, below for qualifications). Top labels and bottom labels cannot be used at the same time.

NOTE: MPC cards do not support the regular hash key configuration. For the MPC-based hash key configuration to be effective, you need an **enhanced-hash-key** configuration.

Load balancing is used to evenly distribute traffic when the following conditions apply:

- There are multiple equal-cost next hops over different interfaces to the same destination.
- There is a single next hop over an aggregated interface.

An LSP tends to load-balance its placement by randomly selecting one of the equal-cost next hops and using it exclusively. The random selection is made independently at each transit router, which compares Interior Gateway Protocol (IGP) metrics alone. No consideration is given to bandwidth or congestion levels.

This feature applies to aggregated Ethernet and aggregated SONET/SDH interfaces as well as multiple equal-cost MPLS next hops. In addition, on the T Series, MX Series, M120, and M320 routers only, you can configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires. You can also configure load balancing for Ethernet pseudowires based on IP information. The option to include IP information in the hash key provides support for Ethernet circuit cross-connect (CCC) connections.

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

```
[edit forwarding-options hash-key]
family mpls {
  all-labels;
  bottom-label-1;
  bottom-label-2;
  bottom-label-3;
  label-1;
  label-2;
  label-3;
  no-labels;
```



```
no-label-1-exp;
payload {
  ether-pseudowire;
  ip {
    disable;
    layer-3-only;
    port-data {
      destination-lsb;
      destination-msb;
      source-lsb;
      source-msb;
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options hash-key]

Table 14 on page 492 provides detailed information about all of the possible MPLS LSP load-balancing options.

Table 14: MPLS LSP Load Balancing Options

Statement	Supported Platforms	MPLS LSP Load Balancing Options
all-labels	MX Series and PTX Series	<p>Prior to Junos OS Release 19.1R1, up to eight MPLS labels were included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. On PTX Series routers, this value is set by default.</p> <p>Starting in Junos OS Release 19.1R1, for MX Series routers with MPC and MIC interfaces, up to sixteen incoming MPLS labels are included in the hash key.</p>
bottom-label1	MX Series with DPC (I-Chip). Not supported on M10i, M7i, and M120.	Uses the bottom-most label for calculating the hash key, for example if the top labels do not provide sufficient variable for the required level of entropy.
bottom-label2	MX Series with DPC (I-Chip). Not supported on M10i, M7i, and M120.	Uses the second label from the bottom for calculating the hash key, for example if the top labels do not provide sufficient variable for the required level of entropy.

Table 14: MPLS LSP Load Balancing Options (*continued*)

Statement	Supported Platforms	MPLS LSP Load Balancing Options
bottom-label-3	MX Series with DPC (I-Chip). Not supported on M10i, M7i, and M120.	Uses the third label from the bottom for calculating the hash key, for example if the top labels do not provide sufficient variable for the required level of entropy.
label-1	M Series, MX Series, T Series	Include the first label in the hash key. Use this option for single label packets.
label-2	M Series, MX Series, T Series	Include the second label in the hash key. You must also configure the label-1 option. The entire first label and the first 16 bits of the second label are used in the hash key.
label-3	M Series, MX Series, T Series	Include the third label in the hash key. You must also configure the label-1 option and the label-2 option.
no-labels	All	Excludes MPLS labels from the hash key.
no-label-1-exp	M Series, MX Series, T Series	Excludes the EXP bit of the top label from the hash key. You must also configure the label-1 option. For Layer 2 VPNs, the router could encounter a packet reordering problem. When a burst of traffic pushes the customer traffic bandwidth to exceed its limits, the traffic might be affected in mid flow. Packets might be reordered as a result. By excluding the EXP bit from the hash key, you can avoid this reordering problem.
payload	All	Allows you to configure which parts of the IP packet payload to include in the hash key. For the PTX Series Packet Transport Router, this value is set by default.
disable	PTX Series	Exclude IP payload from the hash key.
disable-pseudowire	M120, M320, MX Series, T Series	Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
ip	All	Include the IPv4 or IPv6 address in the hash key. You must also configure either label-1 or no-labels .
layer-3-only	All	Include only the Layer 3 IP information in the hash key. Excludes all of the port-data bytes from the hash key.

Table 14: MPLS LSP Load Balancing Options (*continued*)

Statement	Supported Platforms	MPLS LSP Load Balancing Options
port-data	M Series, MX Series, T Series	Include the source and destination port field information. By default, the most significant byte and least significant byte of the source and destination port fields are used in the hash key. To select specific bytes to use in the hash key, include one or more of the source-msb , source-lsb , destination-msb , and destination-lsb options at the [edit forwarding-options hash-key family mpls payload ip port-data] hierarchy level. To prevent all four bytes from being hashed, include the layer-3-only statement at the [edit forwarding-options hash-key family mpls payload ip] hierarchy level.
destination-lsb	M Series, MX Series, T Series	Include the least significant byte of the destination port in the hash key. Can be combined with any of the other port-data options.
destination-msb	M Series, MX Series, T Series	Include the most significant byte of the destination port in the hash key. Can be combined with any of the other port-data options.
source-lsb	M Series, MX Series, T Series	Include the least significant byte of the source port in the hash key. Can be combined with any of the other port-data options.
source-msb	M Series, MX Series, T Series	Include the most significant byte of the source port in the hash key. Can be combined with any of the other port-data options.

The following examples illustrate ways in which you can configure MPLS LSP load balancing:

- To include the IP address as well as the first label in the hash key:
 - For M Series, MX Series, and T Series routers, configure the **label-1** statement and the **ip** option for the **payload** statement at the [edit forwarding-options hash-key family mpls] hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
payload {
  ip;
}
```

- For PTX Series Packet Transport Routers, the **all-labels** and **ip payload** options are configured by default, so no configuration is necessary.
- (M320 and T Series routers only) To include the IP address as well as both the first and second labels in the hash key, configure the **label-1** and **label-2** options and the **ip** option for the **payload** statement at the [edit forwarding-options hash-key family mpls] hierarchy level:


```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
payload {
  ip;
}
```

NOTE: You can include this combination of statements on M320 and T Series routers only. If you include them on an M Series Multiservice Edge Router, only the first MPLS label and the IP payload are used in the hash key.

- For T Series routers, ensure proper load balancing by including the **label-1**, **label-2**, and **label-3** options at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
label-3;
```

- (M Series, MX Series, and T Series routers only) For Layer 2 VPNs, the router could encounter a packet reordering problem. When a burst of traffic pushes the customer traffic bandwidth to exceed its limits, the traffic might be affected in mid flow. Packets might be reordered as a result. By excluding the EXP bit from the hash key, you can avoid this reordering problem. To exclude the EXP bit of the first label from the hash calculations, include the **no-label-1-exp** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
no-label-1-exp;
payload {
  ip;
}
```

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1, for MX Series routers with MPC and MIC interfaces, up to sixteen incoming MPLS labels are included in the hash key.

RELATED DOCUMENTATION

[Configuring Load Balancing for Ethernet Pseudowires | 1308](#)

Configuring Load Balancing Based on MPLS Labels on ACX Series Routers

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

Load balancing is used to evenly distribute traffic when there is a single next hop over an aggregated interface or a LAG bundle. Load balancing using MPLS labels is supported only for LAG interfaces and not for equal-cost multipath (ECMP) links.

By default, when load balancing is used to help distribute traffic, Junos OS employs a hash algorithm to select a next-hop address to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is reselected by means of the hash algorithm. You can configure how the hash algorithm is used to load-balance traffic across interfaces in an aggregated Ethernet (ae) interface.

An LSP tends to load-balance its placement by randomly selecting one of the interfaces in an **ae-** interface bundle and using it exclusively. The random selection is made independently at each transit router, which compares Interior Gateway Protocol (IGP) metrics alone. No consideration is given to bandwidth or congestion levels.

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

```
[edit forwarding-options hash-key]
family mpls {
  all-labels;
  label-1;
  label-2;
  label-3;
  no-labels;
  payload {
    ether-pseudowire;
    ip {
      layer-3-only;
      port-data {
        destination-lsb;
        destination-msb;
        source-lsb;
```



```

        source-msb;
    }
}
}
}

```

You can include this statement at the **[edit forwarding-options hash-key]** hierarchy level.

NOTE: When you configure payload ip (**user@host# set forwarding-options hash-key family mpls payload ip**), configuring **layer-3-only** and **port-data** is mandatory.

Load balancing functionality, without proper hash-keys configuration, may result in an unpredictable behavior.

For Layer 2 VPN/pseudowire tunnel termination, upto two labels are used for hashing and payload MAC destination and source addresses can be optionally selected. These controls can be used to support ether-pseudowire knob in family mpls under hash-key configuration shown above. However, since ACX2000 and ACX4000 also support TDM pseudowires, the ether-pseudowire knobs needs to be used only when TDM pseudowires are not being used.

For Layer 3 VPN tunnel termination, upto two labels are used for hashing and payload IP source and destination addresses and Layer 4 source and destination ports can be optionally selected. These controls can be used for supporting ip port-data knobs in family mpls under hash-key configuration shown above. However, since Layer 4 port MSB and LSB cannot be individually selected, one of destination-lsb or destination-msb knobs or one of source-lsb or source-msb knobs would select Layer 4 destination or source ports, respectively.

For LSR case, upto three labels are used for hashing. If a BOS is seen when parsing the first three labels, BCM examines the first nibble of payload - if the nibble is 4, the payload is treated as IPv4 and if the first nibble is 6, the payload is treated as IPv6 and in such cases payload source and destination IP addresses can be speculatively used for hashing. These controls can be used for supporting ip port-data knobs in family mpls under hash-key configuration. However, Layer 4 ports cannot be used for hashing in LSR case, and only layer-3-only knob is applicable. BCM does not claim support for hashing on fields beyond the three MPLS labels. Load Balancing for a single pseudowire session does not take place in case of LSR as all the traffic specific to that session will carry the same set of MPLS labels.

Load balancing on LSR AE interfaces can be achieved for a higher number of MPLS sessions, that is minimum of 10 sessions. This is applicable for CCC/VPLS/L3VPN. In case of Layer 3 VPN, the traffic may not be equally distributed across the member links as the layer 3 addresses also get accounted for (along with the labels) for the hash input function.

For LER scenarios, in case of ACX5048 and ACX5096, hashing based on Layer 3 and Layer 4 fields is possible by configuring the payload option under the "family mpls" hierarchy. Hashing on the LER is not

be based on Labels. For Layer 3 service, it is mandatory to mention the payload as “layer-3-only” and specify “port-data” in case of Layer 4 service. You can also mention the label count while configuring hash-keys on LER routers.

NOTE: LER and LSR load balancing behavior is applicable for CCC/VPLS/Layer 3 VPN and other IP MPLS scenarios.

This feature applies to aggregated Ethernet and aggregated SONET/SDH interfaces. In addition, you can configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires. You can also configure load balancing for Ethernet pseudowires based on IP information. The option to include IP information in the hash key provides support for Ethernet circuit cross-connect (CCC) connections.

[Table 14 on page 492](#) provides detailed information about all of the possible MPLS LSP load-balancing options.

Table 15: MPLS LSP Load Balancing Options

Statement	MPLS LSP Load Balancing Options
label-1	Include the first label in the hash key. Use this option for single label packets.
label-2	Include the second label in the hash key. You must also configure the label-1 option. The entire first label and the first 16 bits of the second label are used in the hash key.
label-3	Include the third label in the hash key. You must also configure the label-1 option and the label-2 option.
no-labels	Excludes MPLS labels from the hash key.
payload	Allows you to configure which parts of the IP packet payload to include in the hash key. For the PTX Series Packet Transport Switch, this value is set by default.
disable	Exclude IP payload from the hash key.
ether-pseudowire	Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
ip	Include the IPv4 or IPv6 address in the hash key. You must also configure either label-1 or no-labels .
layer-3-only	Include only the Layer 3 IP information in the hash key. Excludes all of the port-data bytes from the hash key.

Table 15: MPLS LSP Load Balancing Options (*continued*)

Statement	MPLS LSP Load Balancing Options
port-data	Include the source and destination port field information. By default, the most significant byte and least significant byte of the source and destination port fields are used in the hash key. To select specific bytes to use in the hash key, include one or more of the source-msb , source-lsb , destination-msb , and destination-lsb options at the [edit forwarding-options hash-key family mpls payload ip port-data] hierarchy level. To prevent all four bytes from being hashed, include the layer-3-only statement at the [edit forwarding-options hash-key family mpls payload ip] hierarchy level.
destination-lsb	Include the least significant byte of the destination port in the hash key. Can be combined with any of the other port-data options.
destination-msb	Include the most significant byte of the destination port in the hash key. Can be combined with any of the other port-data options.
source-lsb	Include the least significant byte of the source port in the hash key. Can be combined with any of the other port-data options.
source-msb	Include the most significant byte of the source port in the hash key. Can be combined with any of the other port-data options.

To include the IP address as well as the first label in the hash key, configure the **label-1** statement and the **ip** option for the **payload** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
payload {
  ip;
}
```

To include the IP address as well as both the first and second labels in the hash key, configure the **label-1** and **label-2** options and the **ip** option for the **payload** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
payload {
  ip;
}
```


Ensure proper load balancing by including the **label-1**, **label-2**, and **label-3** options at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
label-3;
```

RELATED DOCUMENTATION

Configuring Per-Packet Load Balancing

[Configuring Load Balancing for Ethernet Pseudowires](#) | 1308

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

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

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


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

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

```
[edit protocols]
mpls {
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    12.1.1.1 loose;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 10m; # Reserve 10 Mbps
    priority 0 0; # Preemptive, but not preemptable
    adaptive; # Set adaptivity
    primary to-hastings;
    secondary alt-hastings {
      standby;
      bandwidth 1m; # Reserve only 1 Mbps for the secondary path
    }
  }
  interface all;
}
```

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

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

When the LSP is up—either because the primary or secondary path is up, or because both paths are up—the prefix 16.0.0.0/8 is installed in the inet.3 table so that all BGP routes whose BGP next hop falls within that range can use the LSP. Also, the prefix 17.0.0.0/8 is installed in the inet.0 table so that BGP can resolve only its next hop through that prefix. The route also can be reached with the **traceroute** or **ping** command. These two routes are in addition to the 11.1.1.1/32 route.

```
[edit protocols]
mpls {
  admin-groups {
    green 1;
    yellow 2;
    red 3;
  }
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    14.1.1.1 strict;
    13.1.1.1 strict;
    12.1.1.1 strict;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 100m;
    install 16.0.0.0/8; # in inet.3; cannot use to traceroute or ping
    install 17.0.0.0/8 active; # installed in inet.0; can use to traceroute or ping
    primary to-hastings {
      admin-group { # further constraints for path computation
        include-all [ green yellow ];
        exclude red;
      }
      optimize-timer 3600; # reoptimize every hour
    }
    secondary alt-hastings {
      standby;
      no-cspf; # do not perform constrained-path computation
    }
  }
}
interface all;
```


Path Computation for LSPs on an Overloaded Router

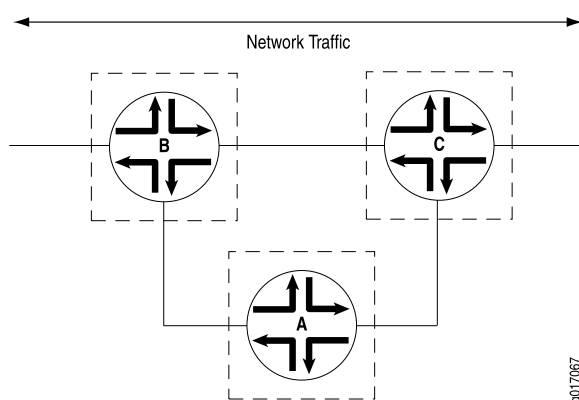
Setting the overload bit in a router running IS-IS causes it to appear overloaded and prevents it from being used for transit traffic. Any new MPLS LSPs, including RSVP-signaled or LDP-signaled LSPs, are re-routed away from an overloaded router. In the case of RSVP, this behavior applies to both Constrained Shortest Path First (CSPF) and non-CSPF LSPs. However, this behavior does not apply to new or existing bypass LSPs. Bypass LSPs are recalculated only when a different event triggers a path recalculation. For example, if you set the smart optimize timer with the [smart-optimize-timer](#) statement, the bypass LSP is rerouted away from the overloaded router only after the specified time elapses. Otherwise, the bypass LSP continues to transit the overloaded router.

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

NOTE: When you set the overload bit on an IS-IS router, any new LSPs transiting through it are recomputed and re-routed away from it. However, existing CSPF LSPs remain active and are not torn down.

An example of when you might want to establish transit LSPs through an overloaded router is illustrated in [Figure 38 on page 503](#), which shows an aggregation router (Router A) dual-homed on two core routers (Router B and Router C). You want to include the aggregation router in the LSP mesh, but transit LSPs should not pass through it, because it is a less capable router with relatively low-bandwidth uplinks to the core. Certain failure and rerouting scenarios could make it impossible for the aggregation router to establish some of its LSPs. Consequently, you run the router in a steady state with the overload bit set, but you are still able to establish ingress and egress LSPs through it.

Figure 38: Aggregation Router A Dual-Homed on Core Routers B and C



RELATED DOCUMENTATION

[Constrained-Path LSP Computation | 485](#)

overload

Computing Backup Paths for LSPs Using Fate Sharing

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

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

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

RELATED DOCUMENTATION

[Configuring the Ingress Router for MPLS-Signaled LSPs | 516](#)

[fate-sharing | 2620](#)

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


Link-state protocols, such as OSPF and IS-IS, use the shortest-path-first (SPF) algorithm to compute the shortest-path tree to all nodes in the network. The results of such computations can be represented by the destination node, next-hop address, and output interface, where the output interface is a physical interface. Label-switched paths (LSPs) can be used to augment the SPF algorithm.

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

The second computation is performed considering only LSPs as a logical interface. Each LSP's egress router is considered. The list of destinations whose shortest path traverses the egress router (established during the first computation) is placed in the inet.3 routing table. These destinations are given the egress router of the LSP as a next hop, enabling BGP on the local router to use these LSPs to access BGP next hops beyond the egress router. Normally, BGP can use only LSPs that terminate at the BGP next hop.


As an illustration, begin with a typical SPF tree (see [Figure 39 on page 505](#)).

Figure 39: Typical SPF Tree, Sourced from Router A



If an LSP connects Router A to Router D and if IGP shortcuts are enabled on Router A, you might have the SPF tree shown in [Figure 40 on page 505](#).

Figure 40: Modified SPF Tree, Using LSP A-D as a Shortcut



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

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

- Use IGP path A-B-D.
- Use LSP A-D.

Router A decides between the two choices by comparing the IGP metrics for path A-B-D with the LSP metrics for LSP A-D. If the IGP metric is lower, path A-B-D is chosen ([Figure 39 on page 505](#)). This path A-B-D is valid only when node D is not the tail-end of the LSP. If node D is the tail end of the LSP, even if the LSP metric is lower or both IGP and LSP metrics are equal, LSP A-D is used ([Figure 40 on page 505](#)).

Note that Router E is reachable through LSP A-D and Router F will take the IGP path.

RELATED DOCUMENTATION

traffic-engineering

OSPF Support for Traffic Engineering

[IGP Shortcuts and Routing Tables](#) | 507

Enabling IGP Shortcuts

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

RELATED DOCUMENTATION

Example: Enabling IS-IS Traffic Engineering Support

Example: Enabling OSPF Traffic Engineering Support

[Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts](#) | 504

LSPs Qualified in IGP Shortcut Computations

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


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

IGP Shortcut Applications

You can use shortcuts to engineer traffic traveling toward destination nodes that do not support MPLS LSPs. For example, in [Figure 40 on page 505](#), traffic traveling toward Router F enters LSP A-E. You can control traffic between Router A and Router F by manipulating LSP A-E; you do not need to explicitly set up an LSP between Router A and Router F.


In [Figure 41 on page 507](#), all traffic from Region 1 to Region 2 traverses LSP A-B if IGP shortcuts are enabled on the ingress router (Router A), permitting aggregation of interregional traffic into one LSP. To perform traffic engineering on the interregional traffic, you have to manipulate LSP A-B only, which avoids creating n2 LSPs from all routers in Region 1 to all routers in Region 2 and allows efficient resource controls on the backbone network.

Figure 41: IGP Shortcuts



Shortcuts allow you to deploy LSPs into a network in an incremental, hierarchical fashion. In [Figure 42 on page 507](#), each region can choose to implement traffic engineering LSPs independently, without requiring cooperation from other regions. Each region can choose to deploy intraregion LSPs to fit the region's bandwidth needs, at the pace appropriate for the region.

Figure 42: IGP Shortcuts in a Bigger Network



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

IGP Shortcuts and Routing Tables

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

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

If traffic engineering for IGP and BGP is enabled (see [“IGP and BGP Destinations” on page 25](#)), IGP moves all routes in inet.3 into inet.0, merging all routes while emptying the inet.3 table. The number of routes in inet.0 will be exactly the same as before. Route next-hops can traverse a physical interface, an LSP, or the combination of the two if the metrics are equal.

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

IGP Shortcuts and VPNs

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

RELATED DOCUMENTATION

[Configuring Traffic Engineering for LSPs | 1137](#)

traffic-engineering


OSPF Support for Traffic Engineering

Advertising LSPs into IGPs

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

As an example, consider the network shown in [Figure 43 on page 508](#).

Figure 43: SPF Computations with Advertised LSPs



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

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

For all routers on the network to know about the LSP between Router E and Router D, you need to advertise it. This advertisement announces the LSP as a unidirectional, point-to-point link in the link-state database, and all routers can compute paths using the LSP. The link-state database maintains information about the AS topology and contains information about the router's local state (for example, the router's usable interfaces and reachable neighbors). In [Figure 43 on page 508](#), Router A will see the link from Router E to Router D and route traffic along this lower-metric path.

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

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

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

Selecting a Forwarding LSP Next Hop

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

RELATED DOCUMENTATION

Configuring the Policy Statement for the Layer 2 Circuit Community

install-nexthop

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

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

```
routing-options {
  router-id 10.10.20.101;
  autonomous-system 2;
  forwarding-table {
    export forwarding-policy;
  }
}
policy-options {
  policy-statement forwarding-policy {
    term one {
      from {
        protocol bgp;
        route-filter 10.1.0.0/16 orlonger;
      }
      then {
        install-nexthop lsp mc-c-lsp-1;
        accept;
      }
    }
    term two {
      from {
        protocol bgp;
        route-filter 10.2.0.0/16 orlonger;
      }
      then {
        install-nexthop lsp mc-c-lsp-2;
        accept;
      }
    }
    term three {
      from {
        protocol bgp;
        route-filter 10.3.0.0/16 orlonger;
      }
      then {
        install-nexthop lsp mc-c-lsp-3;
        accept;
      }
    }
  }
}
```



```

    }
  }
}
protocols {
  mpls {
    label-switched-path mc-c-lsp-1 {
      from 10.10.20.101;
      to 10.10.20.103;
    }
    label-switched-path mc-c-lsp-2 {
      from 10.10.20.101;
      to 10.10.20.103;
    }
    label-switched-path mc-c-lsp-3 {
      from 10.10.20.101;
      to 10.10.20.103;
    }
  }
}

```

ECMP Flow-Based Forwarding on ACX Series Routers

An equal-cost multipath (ECMP) set is formed when the routing table contains multiple next-hop addresses for the same destination with equal cost. (Routes of equal cost have the same preference and metric values.) If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose *one* of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure the Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. On ACX Series routers, per-flow load balancing can be performed to spread traffic across multiple paths between routing devices. ECMP flow-based forwarding is supported for IPv4, IPv6, and MPLS packets on aggregated Ethernet (ae) interfaces.

Load balancing is used to evenly distribute traffic when there are multiple equal-cost next hops over different interfaces or a single next hop over an aggregated interface. By default, when load balancing is used to help distribute traffic, Junos OS employs a hash algorithm to select a next-hop address to install into the forwarding table.

If a next-hop address is no longer part of the ECMP set or if it is removed from the routing table because of a route change, a flow that uses the next hop is rerouted and the session is not affected. Rerouting of the flow also occurs if there is a configuration change that takes away the next-hop address or if an administrator takes down the next-hop interface without deleting it. If a next-hop address is removed

from the routing table because the interface is deleted or the session is intentionally cleared, the session is killed without being rerouted.

To select which packet header data to use for per-flow load balancing, include the **hash-key** statement at the **[edit forwarding-options]** hierarchy level. To load-balance IPv4 traffic by using the port data into the hash key, include the **family-inet** statement at the **[edit forwarding-options hash-key]** hierarchy level. You can incorporate either the Layer 3 IP port data, or the Layer 4 TCP or UDP port data into the hash key. To load-balance based on the MPLS label information, configure the **family mpls** statement at the **[edit forwarding-options hash-key]** hierarchy level.

Forwarding of MPLS traffic by using penultimate-hop popping (PHP) and label-switched routing (LSR) is not supported on ACX Series routers. For ECMP flow-based forwarding over pseudowires, MPLS flows are assigned to one of the ECMP routes by using the hashing algorithm based on user-to-network interface (UNI) index.

To configure ECMP flow-based forwarding on ACX Series routers, first define a load-balancing routing policy by including one or more **policy-statement** configuration statements at the **[edit policy-options]** hierarchy level, with the action **load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** configuration statements at the **[edit routing-options]** hierarchy level.

To view the details of the ECMP next hops and to obtain information for debugging any problem with the ECMP functionality, issue the **show route** or the **show route summary** command.

The maximum number of next-hop addresses in an ECMP set that can be installed in the forwarding table of ACX Series routers is 16. A maximum of 2047 ECMP next-hops are supported.

RELATED DOCUMENTATION

Understanding Routing Policies

Summary of Routing Policy Actions

MPLS LSP Routers

IN THIS CHAPTER

- [Routers in an LSP | 513](#)
- [Configuring the Ingress and Egress Router Addresses for LSPs | 514](#)
- [Configuring the Ingress Router for MPLS-Signaled LSPs | 516](#)
- [Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs | 521](#)
- [Configuring the Connection Between Ingress and Egress Routers | 522](#)
- [Pinging LSPs | 522](#)

Routers in an LSP

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

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

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

Configuring the Ingress and Egress Router Addresses for LSPs

IN THIS SECTION

- [Configuring the Ingress Router Address for LSPs | 514](#)
- [Configuring the Egress Router Address for LSPs | 514](#)
- [Preventing the Addition of Egress Router Addresses to Routing Tables | 515](#)

The following sections describe how to specify the addresses of an LSP's ingress and egress routers:

Configuring the Ingress Router Address for LSPs

The local router always is considered to be the ingress router, which is the beginning of the LSP. The software automatically determines the proper outgoing interface and IP address to use to reach the next router in an LSP.

By default, the router ID is chosen as the address of the ingress router. To override the automatic selection of the source address, specify a source address in the **from** statement:

```
from address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

The outgoing interface used by the LSP is not affected by the source address that you configure.

Configuring the Egress Router Address for LSPs

When configuring an LSP, you must specify the address of the egress router by including the **to** statement:

```
to address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit protocols mpls **static-label-switched-path** *lsp-name*]

- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [static-label-switched-path](#) *lsp-name*]

When you are setting up a signaled LSP, the **to** statement is the only required statement. All other statements are optional.

After the LSP is established, the address of the egress router is installed as a host route in the routing table. This route can then be used by BGP to forward traffic.

To have the software send BGP traffic over an LSP, the address of the egress router is the same as the address of the BGP next hop. You can specify the egress router's address as any one of the router's interface addresses or as the BGP router ID. If you specify a different address, even if the address is on the same router, BGP traffic is not sent over the LSP.

To determine the address of the BGP next hop, use the **show route detail** command. To determine the destination address of an LSP, use the **show mpls lsp** command. To determine whether a route has gone through an LSP, use the **show route** or **show route forwarding-table** command. In the output of these last two commands, the **label-switched-path** or **push** keyword included with the route indicates it has passed through an LSP. Also, use the **traceroute** command to trace the actual path to which the route leads. This is another indication whether a route has passed through an LSP.

You also can manipulate the address of the BGP next hop by defining a BGP import policy filter that sets the route's next-hop address.

Preventing the Addition of Egress Router Addresses to Routing Tables

You must configure an address using the **to** statement for all LSPs. This address is always installed as a /32 prefix in the inet.3 or inet.0 routing tables. You can prevent the egress router address configured using the **to** statement from being added to the inet.3 and inet.0 routing tables by including the **no-install-to-address** statement.

Some reasons not to install the **to** statement address in the inet.3 and inet.0 routing tables include the following:

- Allow Constrained Shortest Path First (CSPF) RSVP LSPs to be mapped to traffic intended for secondary loopback addresses. If you configure an RSVP tunnel, including the **no-install-to-address** statement, and then configure an **install pfx/ <active>** policy later, you can do the following:
 - Verify that the LSP was set up correctly without impacting traffic.
 - Map traffic to the LSP in incremental steps.
 - Map traffic to the destination loopback address (the BGP next hop) by removing the **no-install-to-address** statement once troubleshooting is complete.

- Prevent CCC connections from losing IP traffic. When an LSP determines that it does not belong to a connection, it installs the address specified with the **to** statement in the inet.3 routing table. IP traffic is then forwarded to the CCC remote endpoint, which can cause some types of PICs to fail.

To prevent the egress router address configured using the **to** statement from being added to the inet.3 and inet.0 routing tables, include the **no-install-to-address** statement:

```
no-install-to-address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]
- [edit protocols mpls [static-label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [static-label-switched-path](#) *lsp-name*]

Configuring the Ingress Router for MPLS-Signaled LSPs

IN THIS SECTION

- [Creating Named Paths | 516](#)
- [Configuring Alternate Backup Paths Using Fate Sharing | 518](#)

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

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

Creating Named Paths

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

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

label-switched-path *label-path-name*] hierarchy level. You can specify the same named path on any number of LSPs.

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

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

```
path path-name;
```

You can include this statement at the following hierarchy levels:

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

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

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

You can include this statement at the following hierarchy levels:

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

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

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

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

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

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

Examples: Creating Named Paths

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

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

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

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

Configuring Alternate Backup Paths Using Fate Sharing

IN THIS SECTION

- [Configuring Fate Sharing | 519](#)
- [Implications for CSPF | 520](#)
- [Implications for CSPF When Fate Sharing with Bypass LSPs | 520](#)
- [Example: Configuring Fate Sharing | 521](#)

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

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

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

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

Configuring Fate Sharing

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

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

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

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

Fate-sharing groups contain three types of objects:

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

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


You can list the addresses in any order.

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

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

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

An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.

Implications for CSPF

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

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

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

Implications for CSPF When Fate Sharing with Bypass LSPs

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

- CSPF identifies the fate-sharing groups that are associated with the primary LSP path. CSPF does this by identifying the immediate downstream link and immediate downstream nodes that the bypass is trying to protect. CSPF compiles group lists that contain the immediate downstream link and immediate downstream nodes.

- CSPF checks each link (from ingress to the immediate downstream node) in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group.
- CSPF identifies the downstream link that is not in the fate-shared path.

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

Example: Configuring Fate Sharing

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

```
[edit routing-options]
fate-sharing {
  group east {
    cost 20; # Optional, default value is 1
    from 1.2.3.4 to 1.2.3.5; # A point-to-point link
    from 192.168.200.1; # LAN interface
    from 192.168.200.2; # LAN interface
    from 192.168.200.3; # LAN interface
    from 192.168.200.4; # LAN interface
    from 10.168.1.220; # Router ID of a router node
    from 10.168.1.221; # Router ID of a router node
  }
  group west {
    .....
  }
}
```

Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs

To configure signaled LSPs on all MPLS routers that should participate in MPLS, you need to enable MPLS and RSVP on these routers.

RELATED DOCUMENTATION

[MPLS Configuration Overview](#) | 48

[Minimum RSVP Configuration](#) | 827

Configuring the Connection Between Ingress and Egress Routers

The ingress router might make many attempts to connect and reconnect to the egress router using the primary path. You can control how often the ingress router tries to establish a connection using the primary path and how long it waits between retry attempts.

The retry timer configures how long the ingress router waits before trying to connect again to the egress router using the primary path. The default retry time is 30 seconds. The time can be from 1 through 600 seconds. To modify this value, include the **retry-timer** statement:

```
retry-timer seconds;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

By default, no limit is set to the number of times an ingress router attempts to establish or reestablish a connection to the egress router using the primary path. To limit the number of attempts, include the **retry-limit** statement:

```
retry-limit number;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

The limit can be a value up to 10,000. When the retry limit is exceeded, no more attempts are made to establish a path connection. At this point, intervention is required to restart the primary path.

If you set a retry limit, it is reset to 1 each time a successful primary path is created.

Pinging LSPs

IN THIS SECTION

- [Pinging MPLS LSPs | 523](#)
- [Pinging Point-to-Multipoint LSPs | 523](#)
- [Pinging the Endpoint Address of MPLS LSPs | 524](#)

- [Pinging CCC LSPs | 524](#)
- [Pinging Layer 3 VPNs | 524](#)
- [Support for LSP Ping and Traceroute Commands Based on RFC 4379 | 524](#)

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

Pinging MPLS LSPs

You can ping a specific LSP. Echo requests are sent over the LSP as MPLS packets. The payload is a User Datagram Protocol (UDP) packet forwarded to an address in the 127/8 range (127.0.0.1 by default, this address is configurable) and port 3503. The label and interface information for building and sending this information as an MPLS packet is the same as for standard LSP traffic.

When the echo request arrives at the egress node, the receiver checks the contents of the packet and sends a reply containing the correct return value, by using UDP. The router sending the echo request waits to receive an echo reply after a timeout of 2 seconds (you cannot configure this value).

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router to be able to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

To ping an MPLS LSP use the **ping mpls <count count> <ldp <fec>> <rsvp <exp forwarding-class> <lsp-name>>** command. To ping a secondary MPLS LSP, use the **ping mpls <count count> <rsvp <lsp-name>> standby path-name** command. For a detailed description of this command, see the [CLI Explorer](#).

NOTE: The **ping mpls** command is not supported within routing instances.

NOTE: Self-ping is supported for the master instance and not supported for VLAN-based LSPs or LSPs used in CCC. The message is displayed for each LSP and reduces the readability of the configuration.

Pinging Point-to-Multipoint LSPs

To ping a point-to-multipoint LSP, use the **ping mpls rsvp lsp-name multipoint** or **ping mpls rsvp egress address** commands. The **ping mpls rsvp lsp-name multipoint** command returns a list of all of the egress

router identifiers and the current status of the point-to-multipoint LSP egress routers. The **ping mpls rsvp *lsp-name* multipoint egress *address*** command returns the current status of the specified egress router.

Pinging the Endpoint Address of MPLS LSPs

To determine whether an LSP between two provider edge (PE) routers is up and running, you can ping the endpoint address of the LSP. To ping an MPLS LSP endpoint, use the **ping mpls lsp-end-point *address*** command. This command tells you what type of LSP (RSVP or LDP) terminates at the address specified and whether that LSP is up or down.

For a detailed description of this command, see the [CLI Explorer](#).

Pinging CCC LSPs

You can ping a specific CCC LSP. The CCC LSP ping command is identical to the one used for MPLS LSPs. The command you use is **ping mpls <count *count*> <rsvp <lsp-name>>**. You can also ping a secondary standby CCC LSP by using the **ping mpls <count *count*> <rsvp <lsp-name>> standby *path-name*** command.

For a detailed description of this command, see the [CLI Explorer](#).

Pinging Layer 3 VPNs

You can use a similar command, **ping mpls l3vpn *vpn-name* prefix *prefix* <count *count*>**, to ping a Layer 3 VPN. For more information about this command, see the *Junos OS VPNs Library for Routing Devices* and the [CLI Explorer](#).

Support for LSP Ping and Traceroute Commands Based on RFC 4379

The Junos OS supports LSP **ping** and **traceroute** commands based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*.

LSP **ping** and **traceroute** commands based on RFC 4379 attempt to trace the path taken by an LSP by relying on MPLS TTL expiration. An LSP can take multiple paths from ingress to egress. This occurs in particular with Equal Cost Multipath (ECMP). The LSP **traceroute** command can trace all possible paths to an LSP node.

Configuring MPLS LSPs

IN THIS CHAPTER

- [Configuring LSP Metrics | 526](#)
- [Configuring a Text Description for LSPs | 528](#)
- [Configuring MPLS Soft Preemption | 530](#)
- [Configuring Priority and Preemption for LSPs | 531](#)
- [Configuring Administrative Groups for LSPs | 532](#)
- [Configuring Extended Administrative Groups for LSPs | 534](#)
- [Configuring Preference Values for LSPs | 537](#)
- [Disabling Path Route Recording by LSPs | 537](#)
- [Achieving a Make-Before-Break, Hitless Switchover for LSPs | 538](#)
- [Optimizing Signaled LSPs | 541](#)
- [Configuring the Smart Optimize Timer for LSPs | 545](#)
- [Limiting the Number of Hops in LSPs | 546](#)
- [Configuring the Bandwidth Value for LSPs | 546](#)
- [Automatic Bandwidth Allocation for LSPs | 547](#)
- [Configuring Automatic Bandwidth Allocation for LSPs | 547](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs | 556](#)
- [Configuring an LSP Across ASs | 561](#)
- [Disabling Normal TTL Decrementing | 562](#)
- [Adaptive LSP Configuration | 564](#)
- [Damping Advertisement of LSP State Changes | 565](#)
- [Configuring Primary and Secondary LSPs | 565](#)
- [Configuring Hot Standby of Secondary Paths for LSPs | 568](#)
- [Configuring Corouted Bidirectional LSPs | 570](#)
- [Configuring the Entropy Label for LSPs | 573](#)
- [Example: Configuring an Entropy Label for a BGP Labeled Unicast LSP | 575](#)
- [Configuring Ultimate-Hop Popping for LSPs | 599](#)
- [Configuring Static LSPs | 603](#)
- [Configuring Static Label Switched Paths for MPLS \(CLI Procedure\) | 612](#)

- [Enabling Distributed CSPF for Segment Routing LSPs | 615](#)
- [Color-Based Mapping of VPN Services Overview | 622](#)
- [Configuring Static Label Switched Paths for MPLS | 629](#)
- [Static Segment Routing Label Switched Path | 632](#)
- [Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution | 676](#)
- [Configuring Explicit-Path LSPs | 686](#)
- [Example: Configuring an Explicit-Path LSP | 687](#)
- [Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP | 688](#)

Configuring LSP Metrics

IN THIS SECTION

- [Configuring Dynamic LSP Metrics | 526](#)
- [Configuring Static LSP Metrics | 527](#)

The LSP metric is used to indicate the ease or difficulty of sending traffic over a particular LSP. Lower LSP metric values (lower cost) increase the likelihood of an LSP being used. Conversely, high LSP metric values (higher cost) decrease the likelihood of an LSP being used.

The LSP metric can be specified dynamically by the router or explicitly by the user as described in the following sections:

Configuring Dynamic LSP Metrics

If no specific metric is configured, an LSP attempts to track the IGP metric toward the same destination (the **to** address of the LSP). IGP includes OSPF, IS-IS, Routing Information Protocol (RIP), and static routes. BGP and other RSVP or LDP routes are excluded.

For example, if the OSPF metric toward a router is 20, all LSPs toward that router automatically inherit metric 20. If the OSPF toward a router later changes to a different value, all LSP metrics change accordingly. If there are no IGP routes toward the router, the LSP raises its metric to 65,535.

Note that in this case, the LSP metric is completely determined by IGP; it bears no relationship to the actual path the LSP is currently traversing. If LSP reroutes (such as through reoptimization), its metric does not change, and thus it remains transparent to users. Dynamic metric is the default behavior; no configuration is required.

Configuring Static LSP Metrics

You can manually assign a fixed metric value to an LSP. Once configured with the **metric** statement, the LSP metric is fixed and cannot change:

```
metric number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit protocols mpls **static-label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **static-label-switched-path** *lsp-name*]

The LSP metric has several uses:

- When there are parallel LSPs with the same egress router, the metrics are compared to determine which LSP has the lowest metric value (the lowest cost) and therefore the preferred path to the destination. If the metrics are the same, the traffic is shared.

Adjusting the metric values can force traffic to prefer some LSPs over others, regardless of the underlying IGP metric.

- When an IGP shortcut is enabled (see [“Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts” on page 504](#)), an IGP route might be installed in the routing table with an LSP as the next hop, if the LSP is on the shortest path to the destination. In this case, the LSP metric is added to the other IGP metrics to determine the total path metric. For example, if an LSP whose ingress router is X and egress router is Y is on the shortest path to destination Z, the LSP metric is added to the metric for the IGP route from Y to Z to determine the total cost of the path. If several LSPs are potential next hops, the total metrics of the paths are compared to determine which path is preferred (that is, has the lowest total metric). Or, IGP paths and LSPs leading to the same destination could be compared by means of the metric value to determine which path is preferred.

By adjusting the LSP metric, you can force traffic to prefer LSPs, prefer the IGP path, or share the load among them.

- If router X and Y are BGP peers and if there is an LSP between them, the LSP metric represents the total cost to reach Y from X. If for any reason the LSP reroutes, the underlying path cost might change significantly, but X's cost to reach Y remains the same (the LSP metric), which allows X to report through

a BGP multiple exit discriminator (MED) a stable metric to downstream neighbors. As long as Y remains reachable through the LSP, no changes are visible to downstream BGP neighbors.

It is possible to configure IS-IS to ignore the configured LSP metric by including the **ignore-lsp-metrics** statement at the **[edit protocols isis traffic-engineering shortcuts]** hierarchy level. This statement removes the mutual dependency between IS-IS and MPLS for path computation. For more information, see the *Junos OS Routing Protocols Library*.

Configuring a Text Description for LSPs

You can provide a textual description for an LSP by enclosing any descriptive text that includes spaces within quotation marks (" "). The descriptive text you include is displayed in the detail output of the **show mpls lsp** or the **show mpls container-lsp** command.

Adding a text description for an LSP has no effect on the operation of the LSP. The LSP text description can be no more than 80 characters in length.

To provide a textual description for an LSP, include the **description** statement at any of the following hierarchy levels:

- **[edit protocols mpls label-switched-path lsp-name]**
- **[edit protocols mpls container-label-switched-path lsp-name]**
- **[edit protocols mpls static-label-switched-path lsp-name]**
- **[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name]**
- **[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name]**

Before you begin:

- Configure the device interfaces.
- Configure the device for network communication.
- Enable MPLS on the device interfaces.
- Configure an LSP in the MPLS domain.

To add a text description for an LSP:

1. Enter any text describing the LSP.

```
[edit protocols mpls lsp lsp-name]
user@host# set description text
```

For example:


```
[edit protocols mpls lsp LSP1]
user@host# set description "Connecting remote device"
```

2. Verify and commit the configuration.

For example:

```
[edit protocols mpls lsp]
user@host# set protocols mpls label-switched-path LSP1 to 1.1.1.1
user@host# set protocols mpls label-switched-path LSP1 description "Connecting remote device"
user@host# set protocols mpls interface ge-1/0/8.0
```

```
[edit]
user@host# commit
commit complete
```

3. View the description of an LSP using the **show mpls lsp detail** or **show mpls container-lsp detail** command, depending on the type of LSP configured.

```
user@host> show mpls lsp detail
```

```
Ingress LSP: 1 sessions

1.1.1.1
  From: 0.0.0.0, State: Up, ActiveRoute: 1, LSPname: LSP1
  Description: Connecting remote device
  ActivePath: (none)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    No computed ERO.
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```


RELATED DOCUMENTATION

[label-switched-path](#) | 2653[Configuring LSP Metrics](#) | 526

Configuring MPLS Soft Preemption

Soft preemption attempts to establish a new path for a preempted LSP before tearing down the original LSP. The default behavior is to tear down a preempted LSP first, signal a new path, and then reestablish the LSP over the new path. In the interval between when the path is taken down and the new LSP is established, any traffic attempting to use the LSP is lost. Soft preemption prevents this type of traffic loss. The trade-off is that during the time when an LSP is being soft preempted, two LSPs with their corresponding bandwidth requirements are used until the original path is torn down.

MPLS soft preemption is useful for network maintenance. For example, you can move all LSPs away from a particular interface, then take the interface down for maintenance without interrupting traffic. MPLS soft preemption is described in detail in RFC 5712, *MPLS Traffic Engineering Soft Preemption*.

Soft preemption is a property of the LSP and is disabled by default. You configure it at the ingress of an LSP by including the **soft-preemption** statement:

```
soft-preemption;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]

You can also configure a timer for soft preemption. The timer designates the length of time the router should wait before initiating a hard preemption of the LSP. At the end of the time specified, the LSP is torn down and resigaled. The soft-preemption cleanup timer has a default value of 30 seconds; the range of permissible values is 0 through 180 seconds. A value of 0 means that soft preemption is disabled. The soft-preemption cleanup timer is global for all LSPs.

Configure the timer by including the **cleanup-timer** statement:

```
cleanup-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp [preemption soft-preemption](#)]
- [edit logical-systems *logical-system-name* protocols rsvp [preemption soft-preemption](#)]

NOTE: Soft preemption cannot be configured on LSPs for which fast reroute has been configured. The configuration fails to commit. However, you can enable soft preemption in conjunction with node and link protection.

NOTE: The counter value for *SoftPreemptionCnt* initializes with a value of 0 (zero), visible in the command **show rsvp interface detail** output.

Configuring Priority and Preemption for LSPs

When there is insufficient bandwidth to establish a more important LSP, you might want to tear down a less important existing LSP to free the bandwidth. You do this by preempting the existing LSP.

Whether an LSP can be preempted is determined by two properties associated with the LSP:

- **Setup priority**—Determines whether a new LSP that preempts an existing LSP can be established. For preemption to occur, the setup priority of the new LSP must be higher than that of the existing LSP. Also, the act of preempting the existing LSP must produce sufficient bandwidth to support the new LSP. That is, preemption occurs only if the new LSP can be set up successfully.
- **Reservation priority**—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. When the reservation priority is high, the existing LSP is less likely to give up its reservation, and hence it is unlikely that the LSP can be preempted.

You cannot configure an LSP with a high setup priority and a low reservation priority, because permanent preemption loops might result if two LSPs are allowed to preempt each other. You must configure the reservation priority to be higher than or equal to the setup priority.

The setup priority also defines the relative importance of LSPs on the same ingress router. When the software starts, when a new LSP is established, or during fault recovery, the setup priority determines the order in which LSPs are serviced. Higher-priority LSPs tend to be established first and hence enjoy more optimal path selection.

To configure the LSP's preemption properties, include the **priority** statement:

```
priority setup-priority reservation-priority;
```

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

Both **setup-priority** and **reservation-priority** can be a value from 0 through 7. The value 0 corresponds to the highest priority, and the value 7 to the lowest. By default, an LSP has a setup priority of 7 (that is, it cannot preempt any other LSPs) and a reservation priority of 0 (that is, other LSPs cannot preempt it). These defaults are such that preemption does not happen. When you are configuring these values, the setup priority should always be less than or equal to the hold priority.

Configuring Administrative Groups for LSPs

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups.

Administrative groups are meaningful only when constrained-path LSP computation is enabled.

You can assign up to 32 names and values (in the range 0 through 31), which define a series of names and their corresponding values. The administrative names and values must be identical across all routers within a single domain.

NOTE: The administrative value is distinct from the priority. You configure the priority for an LSP using the **priority** statement. See [“Configuring Priority and Preemption for LSPs” on page 531](#).

To configure administrative groups, follow these steps:

1. Define multiple levels of service quality by including the **admin-groups** statement:

```
admin-groups {
    group-name group-value;
}
```

You can include this statement at the following hierarchy levels:

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

The following configuration example illustrates how you might configure a set of administrative names and values for a domain:

```
[edit protocols mpls]
admin-groups {
    gold 1;
```



```

silver 2;
copper 3;
best-effort 4;
}

```

2. Define the administrative groups to which an interface belongs. You can assign multiple groups to an interface. Include the **interface** statement:

```

interface interface-name {
  admin-group [ group-names ];
}

```

You can include this statement at the following hierarchy levels:

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

If you do not include the **admin-group** statement, an interface does not belong to any group.

IGPs use the group information to build link-state packets, which are then flooded throughout the network, providing information to all nodes in the network. At any router, the IGP topology, as well as administrative groups of all the links, is available.

Changing the interface's administrative group affects only new LSPs. Existing LSPs on the interface are not preempted or recomputed to keep the network stable. If LSPs need to be removed because of a group change, issue the **clear rsvp session** command.

NOTE: When configuring administrative groups and extended administrative groups together for a link, both the types of administrative groups must be configured on the interface.

3. Configure an administrative group constraint for each LSP or for each primary or secondary LSP path. Include the **label-switched-path** statement:

```

label-switched-path lsp-name {
  to address;
  ...
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
}

```



```

primary path-name {
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
}
secondary path-name {
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
}
}

```

You can include the **label-switched-path** statement at the following hierarchy levels:

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

If you omit the **include-all**, **include-any**, or **exclude** statements, the path computation proceeds unchanged. The path computation is based on the constrained-path LSP computation. For information about how the constrained-path LSP computation is calculated, see [“How CSPF Selects a Path” on page 486](#).

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

RELATED DOCUMENTATION

| [Configuring Extended Administrative Groups for LSPs | 534](#)

Configuring Extended Administrative Groups for LSPs

In MPLS traffic engineering, a link can be configured with a set of administrative groups (also known as colors or resource classes). Administrative groups are carried in the interior gateway protocol (IGP) (OSPFv2 and IS-IS) as a 32-bit value assigned to each link. Juniper Networks routers normally interpret this 32-bit

value as a bit mask with each bit representing a group, limiting each network to a total of 32 distinct administrative groups (value range 0 through 31).

You configure extended administrative groups, represented by a 32-bit value, expanding the number of administrative groups supported in the network beyond just 32. The original range of values available for administrative groups is still supported for backwards compatibility.

The extended administrative groups configuration accepts a set of interfaces with a corresponding set of extended administrative group names. It converts the names into a set of 32-bit values and propagates this information into the IGP. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by Constrained Shortest Path First (CSPF) for path computation.

The following procedure describes how to configure extended administrative groups:

1. Configure the **admin-groups-extended-range** statement:

```
admin-groups-extended-range {
    maximum maximum-number;
    minimum minimum-number;
}
```

You can include this statement at the following hierarchy levels:

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

The **admin-groups-extended-range** statement includes the **minimum** and **maximum** options. The range maximum must be greater than the range minimum.

2. Configure the **admin-groups-extended** statement:

```
admin-groups-extended group-name {
    group-value group-identifier;
}
```

You can include this statement at the following hierarchy levels:

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

The **admin-groups-extended** statement enables you to configure a group name and group value for the administrative group. The group value must be within the range of values configured using the **admin-groups-extended-range** statement.

3. The extended administrative groups for an MPLS interface consist of the set of extended administrative group names assigned for the interface. The interface extended administrative group names must be configured for the global extended administrative groups.

To configure an extended administrative group for an MPLS interface, specify the administrative group name within the MPLS interface configuration using the **admin-groups-extended** statement:

```
admin-groups-extended group-name;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls interface *interface-name*]**
 - **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name*]**
4. The LSP extended administrative groups define the set of include and exclude constraints for an LSP and for a path's primary and secondary paths. The extended administrative group names must be configured for the global extended administrative groups.

To configure extended administrative groups for an LSP, include the **admin-group-extended** statement at an LSP hierarchy level:

```
admin-group-extended {
  apply-groups group-value;
  apply-groups-except group-value;
  exclude group-value;
  include-all group-value;
  include-any group-value;
}
```

The **admin-group-extended** statement includes the following options: **apply-groups**, **apply-groups-except**, **exclude**, **include-all**, and **include-any**. Each option enables you to configure one or more extended administrative groups.

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

5. To display the currently configured extended administrative groups, issue the **show mpls admin-groups-extended** command.

NOTE: When configuring administrative groups and extended administrative groups together for a link, both the types of administrative groups must be configured on the interface.

RELATED DOCUMENTATION

[Configuring Administrative Groups for LSPs | 532](#)

Configuring Preference Values for LSPs

As an option, you can configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for RSVP LSPs is 7 and for LDP LSPs is 9. These preference values are lower (more preferred) than all learned routes except direct interface routes.

To change the default preference value, include the **preference** statement:

```
preference preference;
```

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

Disabling Path Route Recording by LSPs

The Junos implementation of RSVP supports the Record Route object, which allows an LSP to actively record the routers through which it transits. You can use this information for troubleshooting and to prevent routing loops. By default, path route information is recorded. To disable recording, include the **no-record** statement:

```
no-record;
```

For a list of hierarchy levels at which you can include the **record** and **no-record** statements, see the statement summary section for the statement.

Achieving a Make-Before-Break, Hitless Switchover for LSPs

IN THIS SECTION

- [Specifying the Amount of Time the Router Waits to Switch Over to New Paths | 539](#)
- [Specifying the Amount of Time to Delay the Tear Down of Old Paths | 539](#)
- [Achieving a Hitless, MBB Switchover Without Artificial Delays | 540](#)

Adaptive label-switched paths (LSPs) might need to establish a new LSP instance and transfer traffic from an old LSP instance onto the new LSP instance before tearing down the old one. This type of configuration is referred to as a *make before break* (MBB).

RSVP-TE is a protocol used to establish LSPs in MPLS networks. The Junos OS implementation of RSVP-TE to achieve a hitless (no traffic loss) MBB switchover has relied on configuring the timer values in the following configuration statements:

- **optimize-switchover-delay**—Amount of time to wait before switching to the new LSP instance.
- **optimize-hold-dead-delay**—Amount of time to wait after switchover and before deletion of the old LSP instance.

Both the **optimize-switchover-delay** and **optimize-hold-dead-delay** statements apply to all LSPs that use the make-before-break behavior for LSP setup and teardown, not just for LSPs for which the **optimize-timer** statement has also been configured. The following MPLS features cause LSPs to be set up and torn down using make-before-break behavior:

- Adaptive LSPs
- Automatic bandwidth allocation
- BFD for LSPs
- Graceful Routing Engine switchover
- Link and node protection
- Nonstop active routing
- Optimized LSPs
- Point-to-multipoint (P2MP) LSPs
- Soft preemption
- Standby secondary paths

Both the **optimize-switchover-delay** and **optimize-hold-dead-delay** statements when configured add an artificial delay to the MBB process. The value of the **optimize-switchover-delay** statement varies with the size of the Explicit Route Objects (EROs). An ERO is an extension to RSVP that allows an RSVP PATH message to traverse an explicit sequence of routers that is independent of conventional shortest-path IP routing. The value of the **optimize-switchover-delay** statement also depends on the CPU load on each of the routers on the path. Customers set the **optimize-switchover-delay** statement by trial and error.

The value of the **optimize-hold-dead-delay** statement depends on how fast the ingress router moves all application prefixes to point to the new LSP. This is determined by the Packet Forwarding Engine load, which can vary from platform to platform. Customers have to set the **optimize-hold-dead-delay** statement by trial and error.

However, as of Release 15.1, Junos OS is able to achieve a hitless MBB switchover without configuring the artificial delays that such timer values introduce.

This topic summarizes the three methods of achieving a MBB switchover from an old LSP to a new LSP using Junos OS:

Specifying the Amount of Time the Router Waits to Switch Over to New Paths

To specify the amount of time the router waits to switch over LSP instances to newly optimized paths, use the **optimize-switchover-delay** statement. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The timer in this statement helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths. This timer can only be enabled or disabled for all of the LSPs configured on the router.

To configure the amount of time the router waits to switch over LSP instances to newly optimized paths, specify the time in seconds by using the **optimize-switchover-delay** statement:

```
optimize-switchover-delay seconds;
```

You can include this statement at the following hierarchy levels:

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

Specifying the Amount of Time to Delay the Tear Down of Old Paths

To specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths, use the **optimize-hold-dead-delay** statement. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The timer in this statement helps to ensure that old paths are not torn down

before all routes have been switched over to the new optimized paths. This timer can be enabled for specific LSPs or for all of the LSPs configured on the router.

To configure the amount of time in seconds to delay the tear down of old paths after the router has switched traffic to new optimized paths, use the **optimize-hold-dead-delay** statement:

```
optimize-hold-dead-delay seconds;
```

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

Achieving a Hitless, MBB Switchover Without Artificial Delays

As of Junos OS Release 15.1, there is another way to relinquish the old LSP instances after MBB switchover without relying on the arbitrary time intervals set up by the **optimize-switchover-delay** or **optimize-hold-dead-delay** statement. For example, if you use the **optimize-hold-dead-delay** statement, you configure a time you think it is safe to wait before tearing down the old LSP instance after MBB. However, some routes might still be in the process of shifting to the new instance. Tearing down the old LSP instance prematurely results in one of the transit nodes dropping the traffic for those routes that have not shifted to the new LSP instance.

To avoid traffic loss, instead of using the **optimize-switchover-delay** statement, you can use MPLS-OAM (lsp ping), which confirms that the LSP data plane is established end-to-end. Instead of using the **optimize-hold-dead-delay** statement, you can use a feedback mechanism from the rpd infrastructure that confirms that all prefixes referring to the old LSP have been switched over. The feedback mechanism is sourced from the Tag library and relies on the routing protocol process (rpd) infrastructure to determine when all the routes using the old LSP instance have fully shifted to the new LSP instance after MBB switchover.

The feedback mechanism is always in place, and it is optional. Configure the **optimize-adaptive-teardown** statement to have the feedback mechanism used during MBB switchover. This feature is not supported for RSVP point-to-multipoint (P2MP) LSP instances. Global configuration of the **optimize-adaptive-teardown** statement only affects the point-to-point LSPs that are configured in the system.

You only need to configure the **optimize-adaptive-teardown** statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). This feedback mechanism ensures that old paths are not torn down before all routes have been switched over to the new optimized paths. The global configuration of this configuration statement affects only the point-to-point LSPs that are configured in the system.

```
optimize-adaptive-teardown {
  p2p:
}
```


You can include this statement at the **[edit protocols mpls]** hierarchy level.

RELATED DOCUMENTATION

[Adaptive LSP Configuration | 564](#)

[Configuring Automatic Bandwidth Allocation for LSPs | 547](#)

[Configuring MPLS Soft Preemption | 530](#)

[Configuring the Smart Optimize Timer for LSPs | 545](#)

[Configuring Hot Standby of Secondary Paths for LSPs | 568](#)

Optimizing Signaled LSPs

Once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A new path might have become available that is less congested, has a lower metric, and traverses fewer hops. You can configure the router to recompute paths periodically to determine whether a more optimal path has become available.

If reoptimization is enabled, an LSP can be rerouted through different paths by constrained-path recomputations. However, if reoptimization is disabled, the LSP has a fixed path and cannot take advantage of newly available network resources. The LSP is fixed until the next topology change breaks the LSP and forces a recomputation.

Reoptimization is not related to failover. A new path is always computed when topology failures occur that disrupt an established path.

Because of the potential system overhead involved, you need to carefully control the frequency of reoptimization. Network stability might suffer when reoptimization is enabled. By default, the **optimize-timer** statement is set to 0 (that is, it is disabled).

LSP optimization is meaningful only when constrained-path LSP computation is enabled, which is the default behavior. For more information about constrained-path LSP computation, see [“Disabling Constrained-Path LSP Computation” on page 489](#). Also, LSP optimization is only applicable to ingress LSPs, so it is only necessary to configure the **optimize-timer** statement on the ingress router. The transit and egress routers require no specific configuration to support LSP optimization (other than to have MPLS enabled).

To enable path reoptimization, include the **optimize-timer** statement:

```
optimize-timer seconds;
```


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

Once you have configured the **optimize-timer** statement, the reoptimization timer continues its countdown to the configured value even if you delete the **optimize-timer** statement from the configuration. The next optimization uses the new value. You can force the Junos OS to use a new value immediately by deleting the old value, committing the configuration, configuring the new value for the **optimize-timer** statement, and then committing the configuration again.

After reoptimization is run, the result is accepted only if it meets the following criteria:

1. The new path is not higher in IGP metric. (The metric for the old path is updated during computation, so if a recent link metric changed somewhere along the old path, it is accounted for.)
2. If the new path has the same IGP metric, it is not more hops away.
3. The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.)
4. The new path does not worsen congestion overall.

The relative congestion of the new path is determined as follows:

- a. The percentage of available bandwidth on each link traversed by the new path is compared to that for the old path, starting from the most congested links.
- b. For each current (old) path, the software stores the four smallest values for bandwidth availability for the links traversed in ascending order.
- c. The software also stores the four smallest bandwidth availability values for the new path, corresponding to the links traversed in ascending order.
- d. If any of the four new available bandwidth values are smaller than any of the corresponding old bandwidth availability values, the new path has at least one link that is more congested than the link used by the old path. Because using the link would cause more congestion, traffic is not switched to this new path.
- e. If none of the four new available bandwidth values is smaller than the corresponding old bandwidth availability values, the new path is less congested than the old path.


When all the above conditions are met, then:

5. If the new path has a lower IGP metric, it is accepted.
6. If the new path has an equal IGP metric and lower hop count, it is accepted.

7. If you choose **least-fill** as a load balancing algorithm, LSPs are load balanced as follows:
- The LSP is moved to a new path that is utilized at least 10% less than the current path. This might reduce congestion on the current path by only a small amount. For example, if an LSP with 1 MB of bandwidth is moved off a path carrying a minimum of 200 MB, congestion on the original path is reduced by less than 1%.
 - For **random** or **most-fill** algorithms, this rule does not apply.

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

Figure 44: least-fill Load Balancing Algorithm Example



As shown in [Figure 44 on page 543](#), there are two potential paths for an LSP to traverse from router A to router H, the odd links from L1 through L13 and the even links from L2 through L14. Currently, the router is using the even links as the active path for the LSP. Each link between the same two routers (for example, router A and router B) has the same bandwidth:

- L1, L2 = 10GE
- L3, L4 = 1GE
- L5, L6 = 1GE
- L7, L8 = 1GE
- L9, L10 = 1GE
- L11, L12 = 10GE
- L13, L14 = 10GE

The 1GE links are more likely to be congested. In this example, the odd 1GE links have the following available bandwidth:

- L3 = 41%
- L5 = 56%
- L7 = 66%
- L9 = 71%

The even 1GE links have the following available bandwidth:

- L4 = 37%
- L6 = 52%

- L8 = 61%
- L10 = 70%

Based on this information, the router would calculate the difference in available bandwidth between the odd links and the even links as follows:

- $L4 - L3 = 41\% - 37\% = 4\%$
- $L6 - L5 = 56\% - 52\% = 4\%$
- $L8 - L7 = 66\% - 61\% = 5\%$
- $L10 - L9 = 71\% - 70\% = 1\%$

The total additional bandwidth available over the odd links is 14% ($4\% + 4\% + 5\% + 1\%$). Since 14% is greater than 10% (the least-fill algorithm minimum threshold), the LSP is moved to the new path over the odd links from the original path using the even links.

8. Otherwise, the new path is rejected.

You can disable the following reoptimization criteria (a subset of the criteria listed previously):

- If the new path has the same IGP metric, it is not more hops away.
- The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.)
- The new path does not worsen congestion overall.
- If the new path has an equal IGP metric and lower hop count, it is accepted.

To disable them, either issue the **clear mpls lsp optimize-aggressive** command or include the **optimize-aggressive** statement:

```
optimize-aggressive;
```

You can include this statement at the following hierarchy levels:

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

Including the **optimize-aggressive** statement in the configuration causes the reoptimization procedure to be triggered more often. Paths are rerouted more frequently. It also limits the reoptimization algorithm to the IGP metric only.

RELATED DOCUMENTATION

[Adaptive LSP Configuration | 564](#)

[Configuring Automatic Bandwidth Allocation for LSPs | 547](#)

[Configuring MPLS Soft Preemption | 530](#)

[Configuring the Smart Optimize Timer for LSPs | 545](#)

[Configuring Hot Standby of Secondary Paths for LSPs | 568](#)

Configuring the Smart Optimize Timer for LSPs

Because of network and router resource constraints, it is typically inadvisable to configure a short interval for the optimize timer. However, under certain circumstances, it might be desirable to reoptimize a path sooner than would normally be provided by the optimize timer.

For example, an LSP is traversing a preferred path that subsequently fails. The LSP is then switched to a less desirable path to reach the same destination. Even if the original path is quickly restored, it could take an excessively long time for the LSP to use it again, because it has to wait for the optimize timer to reoptimize the network paths. For such situations, you might want to configure the smart optimize timer.

When you enable the smart optimize timer, an LSP is switched back to its original path so long as the original path has been restored within 3 minutes of going down. Also, if the original path goes down again within 60 minutes, the smart optimize timer is disabled, and path optimization behaves as it normally does when the optimize timer alone is enabled. This prevents the router from using a flapping link.

The smart optimize timer is dependant on other MPLS features to function properly. For the scenario described here in which an LSP is switched to an alternate path in the event of a failure on the original path, it is assumed that you have configured one or more of the MPLS traffic protection features, including fast reroute, link protection, and standby secondary paths. These features help to ensure that traffic can reach its destination in the event of a failure.

At the least, you must configure a standby secondary path for the smart optimize timer feature to work properly. Fast reroute and link protection are more temporary solutions to a network outage. A secondary path ensures that there is a stable alternate path in the event the primary path fails. If you have not configured any sort of traffic protection for an LSP, the smart optimize timer by itself does not ensure that traffic can reach its destination. For more information about MPLS traffic protection, see [“MPLS and Traffic Protection” on page 476](#).

When a primary path fails and the smart optimize timer switches traffic to the secondary path, the router might continue to use the secondary path even after the primary path has been restored. If the ingress router completes a CSPF calculation, it might determine that the secondary path is the better path.

This might be undesirable if the primary path should be the active path and the secondary path should be used as a backup only. Also, if the secondary path is being used as the active path (even though the primary path has been reestablished) and the secondary path fails, the smart optimize timer feature will not

automatically switch traffic back to the primary path. However, you can enable protection for the secondary path by configuring node and link protection or an additional standby secondary path, in which case, the smart optimize timer can be effective.

Specify the time in seconds for the smart optimize timer using the **smart-optimize-timer** statement:

```
smart-optimize-timer seconds;
```

You can include this statement at the following hierarchy levels:

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

RELATED DOCUMENTATION

[MPLS and Traffic Protection | 476](#)

[Optimizing Signaled LSPs | 541](#)

Limiting the Number of Hops in LSPs

By default, each LSP can traverse a maximum of 255 hops, including the ingress and egress routers. To modify this value, include the **hop-limit** statement:

```
hop-limit number;
```

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

The number of hops can be from 2 through 255. (A path with two hops consists of the ingress and egress routers only.)

Configuring the Bandwidth Value for LSPs

Each LSP has a bandwidth value. This value is included in the sender's Tspec field in RSVP path setup messages. You can specify a bandwidth value in bits per second. If you configure more bandwidth for an LSP, it should be able to carry a greater volume of traffic. The default bandwidth is 0 bits per second.

A nonzero bandwidth requires that transit and egress routers reserve capacity along the outbound links for the path. The RSVP reservation scheme is used to reserve this capacity. Any failure in bandwidth reservation (such as failures at RSVP policy control or admission control) might cause the LSP setup to fail. If there is insufficient bandwidth on the interfaces for the transit or egress routers, the LSP is not established.

To specify a bandwidth value for a signaled LSP, include the **bandwidth** statement:

```
bandwidth bps;
```

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

Automatic Bandwidth Allocation for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth; this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

You set a sampling interval on an LSP configured with automatic bandwidth allocation. The average bandwidth is monitored during this interval. At the end of the interval, an attempt is made to signal a new path for the LSP with the bandwidth allocation set to the maximum average value for the preceding sampling interval. If the new path is successfully established and the original path is removed, the LSP is switched over to the new path. If a new path is not created, the LSP continues to use its current path until the end of the next sampling interval, when another attempt is made to establish a new path. Note that you can set minimum and maximum bandwidth values for the LSP.

During the automatic bandwidth allocation interval, the router might receive a steady increase in traffic (increasing bandwidth utilization) on an LSP, potentially causing congestion or packet loss. To prevent this, you can define a second trigger to prematurely expire the automatic bandwidth adjustment timer before the end of the current adjustment interval.

Configuring Automatic Bandwidth Allocation for LSPs

IN THIS SECTION

- [Configuring Automatic Bandwidth Allocation on LSPs | 549](#)
- [Requesting Automatic Bandwidth Allocation Adjustment | 555](#)

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

At the end of the automatic bandwidth allocation time interval, the current maximum average bandwidth usage is compared with the allocated bandwidth for the LSP. If the LSP needs more bandwidth, an attempt is made to set up a new path where bandwidth is equal to the current maximum average usage. If the attempt is successful, the LSP's traffic is routed through the new path and the old path is removed. If the attempt fails, the LSP continues to use its current path.

NOTE: In calculating the value for **Max AvgBW** (relative to the ingress LSP), the sample collected during make before break (MBB) is ignored to prevent inaccurate results. The first sample after a bandwidth adjustment, or after a change in the LSP ID (regardless of path change), is also ignored.

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:

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

IN THIS SECTION

- [Configuring the Automatic Bandwidth Allocation Interval | 550](#)
- [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth | 550](#)
- [Configuring the Automatic Bandwidth Adjustment Threshold | 551](#)
- [Configuring a Limit on Bandwidth Overflow and Underflow Samples | 552](#)
- [Configuring Passive Bandwidth Utilization Monitoring | 554](#)

To enable automatic bandwidth allocation on an LSP, include the **auto-bandwidth** statement:

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

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 556](#).

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.

NOTE: For a label-switched path (LSP) that has both **bandwidth** and **minimum-bandwidth** for autobandwidth configured under the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level, the LSP bandwidth is adjusted differently.

The LSP is initiated with the bandwidth value configured under the **bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level. At the expiry of the **adjust-interval** timer, the LSP bandwidth gets adjusted based on the traffic flow.

If the bandwidth to be signaled is less than the value configured under the **minimum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the minimum bandwidth.

If the bandwidth to be signaled is greater than the value configured under the **maximum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the maximum bandwidth.

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. Starting in Junos OS Release 14.1R9, 15.1R7, 16.1R5, 16.1X2, 16.2R3, and 17.2R2, all zero value bandwidth samples are considered as underflow samples, except for the zero value samples that arrive after an LSP comes up for the first time, and the zero value samples that arrive first after a Routing Engine switchover.

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 554](#)).
- 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 554](#)), 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.

Release History Table

Release	Description
14.1R9	Starting in Junos OS Release 14.1R9, 15.1R7, 16.1R5, 16.1X2, 16.2R3, and 17.2R2, all zero value bandwidth samples are considered as underflow samples, except for the zero value samples that arrive after an LSP comes up for the first time, and the zero value samples that arrive first after a Routing Engine switchover.

RELATED DOCUMENTATION

[Configuring MPLS to Gather Statistics | 242](#)

[Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs | 556](#)

[request mpls lsp adjust-autobandwidth | 3142](#)

[show mpls lsp | 3214](#)

[show mpls lsp autobandwidth | 3243](#)

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 (MPLS Statistics);
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-transit-statistics;
  transit-statistics-polling;
}
```

2. Specify the **filename** for the files used to store the MPLS trace operation output using the **file** option. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-log**.
3. Specify the maximum number of trace files using the **files number** option. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
4. Specify the interval for calculating the average bandwidth usage by configuring a time in seconds using the **interval** option. You can also set the adjustment interval on a specific LSP by configuring the **interval**

option at the `[edit protocols mpls label-switch-path label-switched-path-name statistics]` hierarchy level.

NOTE: To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (**interval** statement at the `[edit protocols mpls statistics]` hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (**adjust-interval** statement at the `[edit protocols mpls label-switched-path label-switched-path-name auto-bandwidth]` hierarchy level).

5. To trace automatic bandwidth allocation, include the **autobw-state** flag for the MPLS **traceoptions** statement at the `[edit protocols mpls]` hierarchy level.

The following configuration enables the MPLS traceoptions for automatic bandwidth allocation. The trace records are stored in a file called **auto-band-trace** (the filename is user configurable):

```
[edit protocols mpls]
traceoptions {
  file auto-band-trace size 10k files 10 world-readable;
  flag autobw-state;
}
```

6. Using the **show log** command, you can display the automatic bandwidth allocation statistics file generated when you configure the **auto-bandwidth (MPLS Statistics)** 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 1Oct 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 1Oct 30 17:14:57
Total 1 sessions: 1 success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          0 pkt          0 Byte
    0 pps          0 Bps Util 0.00% Reserved Bw          37 Bps
E-D          (LSP ID 6, Tunnel ID 6741)          0 pkt          0 Byte
    0 pps          0 Bps Util 0.00% Reserved Bw          101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh
0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1Oct 30 17:15:27 Total
2 sessions: 2 success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          0 pkt          0 Byte
    0 pps          0 Bps Util 0.00% Reserved Bw          37 Bps
E-D          (LSP ID 6, Tunnel ID 6741)          33 pkt          2695 Byte
    1 pps          89 Bps Util 87.69% Reserved Bw          101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh
0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1Oct 30 17:15:57 Total
2 sessions: 2 success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          0 pkt          0 Byte
    0 pps          0 Bps Util 0.00% Reserved Bw          37 Bps
E-D          (LSP ID 6, Tunnel ID 6741)          65 pkt          5338 Byte
    1 pps          88 Bps Util 86.70% Reserved Bw          101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh
0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1Oct 30 17:16:27 Total
2 sessions: 2 success, 0 fail, 0 ignored
E-D          (LSP ID 6, Tunnel ID 6741)          97 pkt          7981 Byte
    1 pps          88 Bps Util 86.70% Reserved Bw          101 Bps
decr nh 0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1Oct 30 17:16:57
Total 1 sessions: 1 success, 0 fail, 0 ignored

```

7. Issue the **show mpls lsp autobandwidth** command to display current information about automatic bandwidth allocation. The following shows sample output from the **show mpls lsp autobandwidth** command taken at about the same time as the log file shown previously:

```

user@host> show mpls lsp autobandwidth
Lspname          Last          Requested      Reserved      Highwater
AdjustTime LastAdjust
                BW          BW          BW          mark          Left
(sec)

```



```
E-D          300bps      812.005bps  812bps      1.56801kbps 294 sec
Wed Oct 30 17:15:26 2013
```

8. Issue the **file show** command to display the MPLS trace file. You need to specify the file location and file name (the file is located in **/var/log/**). The following shows sample trace file output is taken from an MPLS trace file named **auto-band-trace.0.gz** on a router configured with an LSP named **E-D**. The trace file shows that LSP **E-D** is operating over its reserved bandwidth limit initially. At **Oct 30 17:15:26**, the router triggers an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By **Oct 30 17:15:57**, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```
user@host> file show /var/log/auto-band-trace.0.gz
Oct 30 17:13:57 trace_on: Tracing to "/var/log/E/auto-band-trace" started
Oct 30 17:13:57.466825 LSP E-D (id 5) new bytes arrived          2714 in 29
sec
Oct 30 17:14:27.466713 E-D          (LSP ID 5, Tunnel ID 6741)          241 pkt
          19737 Byte          1 pps          88 Bps Util 234.67% Reserved Bw          37
Bps
Oct 30 17:14:27.466962 LSP E-D (id 5, old id 5); sampled bytes          19737 >
bytes recorded          17094
Oct 30 17:14:27.467035 LSP E-D (id 5) new bytes arrived          2643 in 29
sec
Oct 30 17:14:57.466599 E-D          (LSP ID 5, Tunnel ID 6741)          276 pkt
          22607 Byte          1 pps          95 Bps Util 253.34% Reserved Bw          37
Bps
Oct 30 17:14:57.466758 LSP E-D (id 5, old id 5); sampled bytes          22607 >
bytes recorded          19737
Oct 30 17:14:57.466825 LSP E-D (id 5) new bytes arrived          2870 in 29
sec
Oct 30 17:15:26.265816 Adjust Autobw: LSP E-D (id 5) curr adj bw 300bps updated
with 812.005bps
Oct 30 17:15:26.266064 mpls LSP E-D Autobw change 512.005bps >= threshold 75bps
Oct 30 17:15:26.363372 Autobw Success: LSP E-D () (old id 5 new id 6) update
prev active bw 300 bps with 812 bps
Oct 30 17:15:26.363686 RPD_MPLS_PATH_BANDWIDTH_CHANGE: MPLS path (lsp E-D)
bandwidth changed, path bandwidth 812 bps
Oct 30 17:15:27.364751 RPD_MPLS_LSP_BANDWIDTH_CHANGE: MPLS LSP E-D bandwidth
changed, lsp bandwidth 812 bps
Oct 30 17:15:27.466849 E-D          (LSP ID 5, Tunnel ID 6741)          0 pkt
          0 Byte          0 pps          0 Bps Util 0.00% Reserved Bw          37
Bps
Oct 30 17:15:27.467050 E-D          (LSP ID 6, Tunnel ID 6741)          0 pkt
```



```

0 Byte      0 pps      0 Bps Util  0.00% Reserved Bw      101
Bps
Oct 30 17:15:57.466858 E-D      (LSP ID 5, Tunnel ID 6741)      0 pkt
0 Byte      0 pps      0 Bps Util  0.00% Reserved Bw      37
Bps
Oct 30 17:15:57.467106 E-D      (LSP ID 6, Tunnel ID 6741)      33 pkt
2695 Byte    1 pps      89 Bps Util 87.69% Reserved Bw      101
Bps
Oct 30 17:15:57.467201 LSP E-D (id 6, old id 5); LSP up after autobw adjustment
and active for 30 sec
Oct 30 17:15:57.467398 LSP E-D (id 6) psb bytes      2695 < bytes recorded
22607 total bytes      2695 in 30 sec
Oct 30 17:15:57.467461 First sample of the adjust interval after automatic bw
adjustment
Oct 30 17:15:57.467594 Update curr max avg bw 0bps of LSP E-D with new bw
716.225bps
Oct 30 17:16:27.466830 E-D      (LSP ID 5, Tunnel ID 6741)      0 pkt
0 Byte      0 pps      0 Bps Util  0.00% Reserved Bw      37
Bps
Oct 30 17:16:27.467079 E-D      (LSP ID 6, Tunnel ID 6741)      65 pkt
5338 Byte    1 pps      88 Bps Util 86.70% Reserved Bw      101
Bps
Oct 30 17:16:27.467171 LSP E-D (id 6, old id 6); sampled bytes      5338 >
bytes recorded      2695
Oct 30 17:16:27.467237 LSP E-D (id 6) new bytes arrived      2643 in 29
sec
Oct 30 17:16:57.466712 E-D      (LSP ID 6, Tunnel ID 6741)      97 pkt
7981 Byte    1 pps      88 Bps Util 86.70% Reserved Bw      101
Bps
Oct 30 17:16:57.466870 LSP E-D (id 6, old id 6); sampled bytes      7981 >
bytes recorded      5338

```

RELATED DOCUMENTATION

Configuring Automatic Bandwidth Allocation for LSPs 547
show mpls lsp autobandwidth 3243

Configuring an LSP Across ASs

You can configure an LSP to traverse multiple areas in a network by including the **inter-domain** statement as a part of the LSP configuration. This statement allows the router to search for routes in the IGP database. You need to configure this statement on routers that might be unable to locate a path using intra-domain CSPF (by looking in the traffic engineering database (TED)). When you configure inter-area LSPs, the **inter-domain** statement is required.

Before you begin:

- Configure the device interfaces with family MPLS.
- Configure the device router ID and autonomous system number.
- Enable MPLS and RSVP on the router and transit interfaces.
- Configure your IGP to support traffic engineering.
- Set up an LSP from the ingress to the egress router.

To configure an LSP across multiple ASs on the ingress label-switched router (LER):

1. Enable MPLS on all the interfaces (excluding the management interface).

```
[edit protocols]
user@LER# set mpls interface all
user@LER# set mpls interface fxp0.0 disable
```

2. Enable RSVP on all the interfaces (excluding the management interface).

```
[edit protocols]
user@LER# set rsvp interface all
user@LER# set rsvp interface fxp0.0 disable
```

3. Configure the inter-area LSP.

```
[edit protocols]
user@LER# set mpls label-switched-path inter-area-LSP-name to egress-LER-ip-address
user@LER# set mpls label-switched-path inter-area-LSP-name inter-domain
```

4. Verify and commit the configuration.

```
[edit protocols]
user@LER# set rsvp interface ge-0/0/0.0
```



```

user@LER# set rsvp interface lo0.0
user@LER# set rsvp interface fxp0.0 disable
user@LER# set mpls statistics traffic-class-statistics
user@LER# set mpls label-switched-path R1-R2 to 20.0.0.1
user@LER# set mpls label-switched-path R1-R2 inter-domain
user@LER# set mpls interface ge-0/0/0.0
user@LER# set mpls interface lo0.0
user@LER# set mpls interface fxp0.0 disable
user@LER# set ospf traffic-engineering
user@LER# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@LER# set ospf area 0.0.0.0 interface lo0.0

```

RELATED DOCUMENTATION

[inter-domain](#) | [2650](#)

Disabling Normal TTL Decrementing

By default, the time-to-live (TTL) field value in the packet header is decremented by 1 for every hop the packet traverses in the LSP, thereby preventing loops. If the TTL field value reaches 0, packets are dropped, and an Internet Control Message Protocol (ICMP) error packet is sent to the originating router.

If the normal TTL decrement is disabled, the TTL field of IP packets entering LSPs are decremented by only 1 on transiting the LSP, making the LSP appear as a one-hop router to diagnostic tools, such as **traceroute**. Decrementing the TTL field by 1 is done by the ingress router, which pushes a label on IP packets with the TTL field in the label initialized to 255. The label's TTL field value is decremented by 1 for every hop the MPLS packet traverses in the LSP. On the penultimate hop of the LSP, the router pops the label but does not write the label's TTL field value to the IP packet's TTL field. Instead, when the IP packet reaches the egress router, the IP packet's TTL field value is decremented by 1.

When you use **traceroute** to diagnose problems with an LSP from outside that LSP, **traceroute** sees the ingress router, even though the egress router performs the TTL decrement. The behavior of **traceroute** is different if it is initiated from the ingress router of the LSP. In this case, the egress router would be the first router to respond to **traceroute**.

You can disable normal TTL decrementing in an LSP so that the TTL field value does not reach 0 before the packet reaches its destination, thus preventing the packet from being dropped. You can also disable normal TTL decrementing to make the MPLS cloud appear as a single hop, thereby hiding the network topology.

There are two ways to disable TTL decrementing:

- On the ingress of the LSP, if you include the **no-decrement-ttl** statement, the ingress router negotiates with all downstream routers using a proprietary RSVP object, to ensure all routers are in agreement. If negotiation succeeds, the whole LSP behaves as one hop to transit IP traffic.

```
no-decrement-ttl;
```

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

NOTE: The RSVP object is proprietary to the Junos OS and might not work with other software. This potential incompatibility applies only to RSVP-signaled LSPs. When you include the **no-decrement-ttl** statement, TTL hiding can be enforced on a per-LSP basis.

- On the ingress router, you can include the **no-propagate-ttl** statement. The **no-propagate-ttl** statement applies to all LSPs, regardless of whether they are RSVP-signaled or LDP-signaled. Once set, all future LSPs traversing through this router behave as a single hop to IP packets. LSPs established before you configure this statement are not affected.

```
no-propagate-ttl;
```

You can include this statement at the following hierarchy levels:

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

The operation of the **no-propagate-ttl** statement is interoperable with other vendors' equipment. However, you must ensure that all routers are configured identically.

To configure the TTL behavior for a single VRF routing instance, include the **no-vrf-propagate-ttl** or the **vrf-propagate-ttl** statement in the routing instance configuration at the [edit routing-instances *instance-name*] hierarchy level. The **no-vrf-propagate-ttl** or the **vrf-propagate-ttl** statement overrides the behavior configured globally for the router. If the router is operating in default mode with normal TTL decrementing, the **no-vrf-propagate-ttl** overrides the global behavior for the routing instance on which the **no-vrf-propagate-ttl** statement is configured.

RELATED DOCUMENTATION

Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing (on Layer 3 VPNs User Guide for Routing Devices in the Junos VPNs Configuration Guide)

Adaptive LSP Configuration

An LSP occasionally might need to reroute itself for these reasons:

- The continuous reoptimization process is configured with the **optimize-timer** statement.
- The current path has connectivity problems.
- The LSP is preempted by another LSP configured with the **priority** statement and is forced to reroute.
- The explicit-path information for an active LSP is modified, or the LSP's bandwidth is increased.

You can configure an LSP to be *adaptive* when it is attempting to reroute itself. When it is adaptive, the LSP holds onto existing resources until the new path is successfully established and traffic has been cut over to the new LSP. To retain its resources, an adaptive LSP does the following:

- Maintains existing paths and allocated bandwidths—This ensures that the existing path is not torn down prematurely and allows the current traffic to continue flowing while the new path is being set up.
- Avoids double-counting for links that share the new and old paths—Double-counting occurs when an intermediate router does not recognize that the new and old paths belong to the same LSP and counts them as two separate LSPs, requiring separate bandwidth allocations. If some links are close to saturation, double-counting might cause the setup of the new path to fail.

By default, adaptive behavior is disabled. You can include the **adaptive** statement in two different hierarchy levels.

If you specify the **adaptive** statement at the LSP hierarchy levels, the adaptive behavior is enabled on all primary/secondary paths of the LSP. This means both the primary and secondary paths share the same bandwidth on common links.

To configure adaptive behavior for all LSP paths, include the **adaptive** statement in the LSP configuration:

```
adaptive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

If you specify the **adaptive** statement at the [edit protocols mpls **label-switched-path** *lsp-name* (primary | secondary) *path-name*] hierarchy level, adaptive behavior is enabled only on the path on which it is specified. Bandwidth double-counting occurs between different paths. However, if you also have the **adaptive** statement configured at the [edit protocols mpls **label-switched-path** *lsp-name*] hierarchy level, it overrides the adaptive behavior of each individual path.

To configure adaptive behavior for either the primary or secondary level, include the **adaptive** statement:


```
adaptive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* (**primary** | **secondary**) *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* (**primary** | **secondary**) *path-name*]

Damping Advertisement of LSP State Changes

When an LSP changes from being up to being down, or from down to up, this transition takes effect immediately in the router software and hardware. However, when advertising LSPs into IS-IS and OSPF, you may want to damp LSP transitions, thereby not advertising the transition until a certain period of time has transpired (known as the hold time). In this case, if the LSP goes from up to down, the LSP is not advertised as being down until it has remained down for the hold-time period. Transitions from down to up are advertised into IS-IS and OSPF immediately. Note that LSP damping affects only the IS-IS and OSPF advertisements of the LSP; other routing software and hardware react immediately to LSP transitions.

To damp LSP transitions, include the **advertisement-hold-time** statement:

```
advertisement-hold-time seconds;
```

seconds can be a value from 0 through 65,535 seconds. The default is 5 seconds.

You can include this statement at the following hierarchy levels:

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

Configuring Primary and Secondary LSPs

IN THIS SECTION

- [Configuring Primary and Secondary Paths for an LSP | 566](#)
- [Configuring the Revert Timer for LSPs | 567](#)
- [Specifying the Conditions for Path Selection | 567](#)

By default, an LSP routes itself hop-by-hop toward the egress router. The LSP tends to follow the shortest path as dictated by the local routing table, usually taking the same path as destination-based, best-effort traffic. These paths are “soft” in nature because they automatically re-route themselves whenever a change occurs in a routing table or in the status of a node or link.

To configure the path so that it follows a particular route, create a named path using the **path** statement, as described in [“Creating Named Paths” on page 516](#). Then apply the named path by including the **primary** or **secondary** statement. A named path can be referenced by any number of LSPs.

To configure primary and secondary paths for an LSP, complete the steps in the following sections:

Configuring Primary and Secondary Paths for an LSP

The **primary** statement creates the primary path, which is the LSP’s preferred path. The **secondary** statement creates an alternative path. If the primary path can no longer reach the egress router, the alternative path is used.

To configure primary and secondary paths, include the **primary** and **secondary** statements:

```
primary path-name {
    ...
}
secondary path-name {
    ...
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls [label-switched-path lsp-name](#)]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path lsp-name](#)]

When the software switches from the primary to a secondary path, it continuously attempts to revert to the primary path, switching back to it when it is again reachable, but no sooner than the retry time specified in the **retry-timer** statement. (For more information, see [“Configuring the Connection Between Ingress and Egress Routers” on page 522](#).)

You can configure zero or one primary path. If you do not configure a primary path, the first secondary path that is established is selected as the path.

You can configure zero or more secondary paths. All secondary paths are equal. The software does not attempt to switch among secondary paths. If the current secondary path is not available, the next one is tried in no particular order. To create a set of equal paths, specify secondary paths without specifying a primary path.

If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary to reach the egress router.

Configuring the Revert Timer for LSPs

For LSPs configured with both primary and secondary paths, it is possible to configure the revert timer. If a primary path goes down and traffic is switched to the secondary path, the revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert traffic back to a primary path. If during this time, the primary path experiences any connectivity problems or stability problems, the timer is restarted. You can configure the revert timer for both static and dynamic LSPs.

The Junos OS also makes a determination as to which path is the preferred path. The preferred path is the path that has not encountered any difficulty in the last revert timer period. If both the primary and secondary paths have encountered difficulty, neither path is considered preferred. However, if one of the paths is dynamic and the other static, the dynamic path is selected as the preferred path.

If you have configured BFD on the LSP, Junos OS waits until the BFD session comes up on the primary path before starting the revert timer counter.

The range of values you can configure for the revert timer is 0 through 65,535 seconds. The default value is 60 seconds.

If you configure a value of 0 seconds, the traffic on the LSP, once switched from the primary path to the secondary path, remains on the secondary path permanently (until the network operator intervenes or until the secondary path goes down).

You can configure the revert timer for all LSPs on the router at the `[edit protocols mpls]` hierarchy level or for a specific LSP at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level.

To configure the revert timer, include the **revert-timer** statement:

```
revert-timer seconds;
```

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

Specifying the Conditions for Path Selection

When you have configured both primary and secondary paths for an LSP, you may need to ensure that only a specific path is used.

The **select** statement is optional. If you do not include it, MPLS uses an automatic path selection algorithm.

The **manual** and **unconditional** options do the following:

- **manual**—The path is immediately selected for carrying traffic as long as it is up and stable. Traffic is sent to other working paths if the current path is down or degraded (receiving errors). This parameter overrides all other path attributes except the **select unconditional** statement.

- **unconditional**—The path is selected for carrying traffic unconditionally, regardless of whether the path is currently down or degraded (receiving errors). This parameter overrides all other path attributes.

Because the **unconditional** option switches to a path without regard to its current status, be aware of the following potential consequences of specifying it:

- If a path is not currently up when you enable the **unconditional** option, traffic can be disrupted. Ensure that the path is functional before specifying the **unconditional** option.
- Once a path is selected because it has the **unconditional** option enabled, all other paths for the LSP are gradually cleared, including the primary and standby paths. No path can act as a standby to an unconditional path, so signaling those paths serves no purpose.

For a specific path, the **manual** and **unconditional** options are mutually exclusive. You can include the **select** statement with the **manual** option in the configuration of only one of an LSP's paths, and the **select** statement with the **unconditional** option in the configuration of only one other of its paths.

Enabling or disabling the **manual** and **unconditional** options for the **select** statement while LSPs and their paths are up does not disrupt traffic.

To specify that a path be selected for carrying traffic if it is up and stable for at least the revert timer window, include the **select** statement with the **manual** option:

```
select manual;
```

To specify that a path should always be selected for carrying traffic, even if it is currently down or degraded, include the **select** statement with the **unconditional** option:

```
select unconditional;
```

You can include the **select** statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* (**primary** | **secondary**) *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* (**primary** | **secondary**) *path-name*]

Configuring Hot Standby of Secondary Paths for LSPs

By default, secondary paths are set up only as needed. To have the system maintain a secondary path in a hot-standby state indefinitely, include the **standby** statement:

```
standby;
```


You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **secondary**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **secondary**]

The hot-standby state is meaningful only on secondary paths. Maintaining a path in a hot-standby state enables swift cutover to the secondary path when downstream routers on the current active path indicate connectivity problems. Although it is possible to configure the **standby** statement at the [edit protocols mpls **label-switched-path** *lsp-name* **primary path-name**] hierarchy level, it has no effect on router behavior.

If you configure the **standby** statement at the following hierarchy levels, the hot-standby state is activated on all secondary paths configured beneath that hierarchy level:

- [edit protocols mpls]
- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

The hot-standby state has two advantages:

- It eliminates the call-setup delay during network topology changes. Call setup can suffer from significant delays when network failures trigger large numbers of LSP reroutes at the same time.
- A cutover to the secondary path can be made before RSVP learns that an LSP is down. There can be significant delays between the time the first failure is detected by protocol machinery (which can be an interface down, a neighbor becoming unreachable, a route becoming unreachable, or a transient routing loop being detected) and the time an LSP actually fails (which requires a timeout of soft state information between adjacent RSVP routers). When topology failures occur, hot-standby secondary paths can usually achieve the smallest cutover delays with minimal disruptions to user traffic.

When the primary path is considered to be stable again, traffic is automatically switched from the standby secondary path back to the primary path. The switch is performed no faster than twice the retry-timer interval and only if the primary path exhibits stability throughout the entire switch interval.

The drawback of the hot-standby state is that more state information must be maintained by all the routers along the path, which requires overhead from each of the routers.


NOTE: When viewed with **inet.3**, the same LSP may appear to be shown twice as the active route (both primary and secondary), even though traffic actually is being forwarded over the primary path LSP only. This is normal output, and reflects only that the secondary standby path is available.

Configuring Corouted Bidirectional LSPs

A corouted bidirectional packet LSP is a combination of two LSPs sharing the same path between a pair of ingress and egress nodes, as shown in [Figure 45 on page 570](#). It is established using the GMPLS extensions to RSVP-TE. This type of LSP can be used to carry any of the standard types of MPLS-based traffic, including Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs. You can configure a single BFD session for the bidirectional LSP (you do not need to configure a BFD session for each LSP in each direction). You can also configure a single standby bidirectional LSP to provide a backup for the primary bidirectional LSP. Corouted bidirectional LSPs are supported for both penultimate hop popping (PHP) and ultimate hop popping (UHP).

High availability is available for bidirectional LSPs. You can enable graceful restart and nonstop active routing. Graceful restart and nonstop active routing are supported when the restarting router is the ingress, egress, or transit router for the bidirectional LSP.

Figure 45: Corouted Bidirectional LSP



To configure a corouted bidirectional LSP:

1. In configuration mode, configure the ingress router for the LSP and include the **corouted-bidirectional** statement to specify that the LSP be established as a corouted bidirectional LSP.

The path is computed using CSPF and initiated using RSVP signaling (just like a unidirectional RSVP signaled LSP). Both the path to the egress router and the reverse path from the egress router are created when this configuration is committed.

```
[edit protocols mpls]
user@PE1# set label-switched-path sample-lsp corouted-bidirectional
```

2. (Optional) For a reverse path, configure an LSP on the egress router and include the **corouted-bidirectional-passive** statement to associate the LSP with another LSP.

No path computation or signaling is used for this LSP since it relies on the path computation and signaling provided by the ingress LSP. You cannot configure both the **corouted-bidirectional** statement and the **corouted-bidirectional-passive** statement on the same LSP.


```
[edit protocols mpls]
user@PE1# set label-switched-path sample-lsp-reverse-path corouted-bidirectional-passive
```

This statement also makes it easier to debug corouted bidirectional LSPs. If you configure the **corouted-bidirectional-passive** statement (again, on the egress router), you can issue **ping mpls lsp-end-point**, **ping mpls ldp**, **ping mpls rsvp**, **traceroute mpls ldp**, and **traceroute mpls rsvp** commands to test the corouted bidirectional LSP from the egress router.

3. Use the **show mpls lsp extensive** and the **show rsvp session extensive** commands to display information about the bidirectional LSP.

The following shows output for the **show rsvp session extensive** command when run on an ingress router with a bidirectional LSP configured:

```
user@PE1> show rsvp session extensive
```

```
Ingress RSVP: 2 sessions

10.255.14.39
  From: 10.255.14.43, LSPstate: Up, ActiveRoute: 0
  LSPname: l-to-h, LSPpath: Primary
  LSPtype: Static Configured
  Bidirectional, Upstream label in: 3, Upstream label out: -
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300032
  Resv style: 1 FF, Label in: -, Label out: 300032
  Time left: -, Since: Tue May 31 08:49:25 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 24617 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.1.1.2 (ge-0/0/0.0) 3396 pkts
  RESV rcvfrom: 10.1.1.2 (ge-0/0/0.0) 3394 pkts
  PATH notifyto: localclient
  RESV notifyto: 10.255.14.39
  Protection attributes: primary, working, 1:N protection
  Association attributes: recovery, src 10.255.14.43, id 1
  Explot route: 10.1.1.2 10.1.2.2 10.1.3.2
  Record route: 10.1.1.2 10.1.2.2 10.1.3.2

10.255.14.39
  From: 10.255.14.43, LSPstate: Up, ActiveRoute: 0
  LSPname: l-to-h, LSPpath: Secondary
```



```

LSPtype: Static Configured
Bidirectional, Upstream label in: 3, Upstream label out: -
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 300032
Resv style: 1 FF, Label in: -, Label out: 300032
Time left:    -, Since: Tue May 31 08:49:25 2011
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 24617 protocol 0
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.1.1.2 (ge-0/0/0.0) 3396 pkts
RESV rcvfrom: 10.1.1.2 (ge-0/0/0.0) 3394 pkts
Protection attributes: primary, protecting
Association attributes: recovery, src 10.255.14.43, id 1
Explct route: 10.2.1.2 10.2.2.2 10.2.3.2
Record route: 10.2.1.2 10.2.2.2 10.2.3.2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

RELATED DOCUMENTATION

[Configuring Ultimate-Hop Popping for LSPs | 599](#)

[Configuring LDP Graceful Restart | 960](#)

[Configuring RSVP Graceful Restart | 883](#)

[Configuring Nonstop Active Routing](#)

Configuring the Entropy Label for LSPs

The insertion of entropy labels for an LSP enables transit routers to load-balance MPLS traffic across ECMP paths or Link Aggregation groups using just the MPLS label stack as a hash input without having to rely on deep packet inspection. Deep packet inspection requires more of the router's processing power and different routers have differing deep-packet inspection capabilities.

To configure the entropy label for an LSP, complete the following steps:

1. On the ingress router, include the **entropy-label** statement at the **[edit protocols mpls labeled-switched-path *labeled-switched-path-name*]** hierarchy level or at the **[edit protocols mpls static-labeled-switched-path *labeled-switched-path-name* ingress]** hierarchy level. The entropy label is added to the MPLS label stack and can be processed in the forwarding plane.

```
entropy-label;
```

NOTE: This is only applicable for RSVP and static LSPs.

2. On the ingress router, you can configure an ingress policy for LDP-signaled LSPs:

```
entropy-label {
  ingress-policy policy-name;
}
```

Configure the ingress policy at the **[edit policy-options]** hierarchy level:

```
policy-options {
  policy-statement policy-name {
    term term-name {
      from {
        prefix-list prefix-list-name;
      }
      then actions;
    }
  }
}
```


The following shows an example of an entropy label ingress policy.

```
policy-options {  
  policy-statement entropy-policy {  
    term no-insert-entropy-label {  
      from {  
        prefix-list no-entropy-label-fec;  
      }  
      then accept;  
    }  
  }  
}
```

3. (Optional) By default, routers that support the pushing and popping of entropy labels are configured with the **load-balance-label-capability** statement at the **[edit forwarding-options]** hierarchy level to signal the labels on a per-LSP basis. If the peer router is not equipped to handle load-balancing labels, you can prevent the provider edge (PE) router from signaling the entropy label capability by configuring the **no-load-balance-label-capability** statement at the **[edit forwarding-options]** hierarchy level.

```
[edit forwarding-options]  
user@PE no-load-balance-label-capability;
```

Transit routers require no configuration. The presence of the entropy label indicates to the transit router to load balance based solely on the MPLS label stack.

Penultimate hop routers pop the entropy label by default.

RELATED DOCUMENTATION

[entropy-label](#) | 2602

[ingress-policy](#) | 2647

Example: Configuring an Entropy Label for a BGP Labeled Unicast LSP

IN THIS SECTION

- [Requirements | 575](#)
- [Overview | 576](#)
- [Configuration | 577](#)
- [Verification | 595](#)

This example shows how to configure an entropy label for a BGP labeled unicast to achieve end-to-end load balancing using entropy labels. When an IP packet has multiple paths to reach its destination, Junos OS uses certain fields of the packet headers to hash the packet to a deterministic path. This requires an entropy label, a special load-balancing label that can carry the flow information. LSRs in the core simply use the entropy label as the key to hash the packet to the correct path. An entropy label can be any label value between 16 to 1048575 (regular 20-bit label range). Since this range overlaps with the existing regular label range, a special label called entropy label indicator (ELI) is inserted before the entropy label. ELI is a special label assigned by IANA with the value of 7.

BGP labeled unicasts generally concatenate RSVP or LDP LSPs across multiple IGP areas or multiple autonomous systems. RSVP or LDP entropy labels are popped at the penultimate hop node, together with the RSVP or LDP label. This feature enables the use of entropy labels at the stitching points to bridge the gap between the penultimate hop node and the stitching point, in order to achieve end-to-end entropy label load balancing for BGP traffic.

Requirements

This example uses the following hardware and software components:

- Seven MX Series routers with MPCs
- Junos OS Release 15.1 or later running on all the devices

Before you configure an entropy label for BGP labeled unicast, make sure you:

1. Configure the device interfaces.
2. Configure OSPF or any other IGP protocol.
3. Configure BGP.

4. Configure RSVP.
5. Configure MPLS.

Overview

When BGP labeled unicasts concatenate RSVP or LDP LSPs across multiple IGP areas or multiple autonomous systems, RSVP or LDP entropy labels are popped at the penultimate hop node, together with the RSVP or LDP label. However, there are no entropy labels at the stitching points, that is, the routers between two areas. Therefore, the routers at the stitching points used the BGP labels to forward packets.

Beginning with Junos OS Release 15.1, you can configure an entropy label for BGP labeled unicast to achieve end-to-end entropy label load balancing. This feature enables the use of an entropy label at the stitching points in order to achieve end-to-end entropy label load balancing for BGP traffic. Junos OS allows the insertion of entropy labels at the BGP labeled unicast LSP ingress.

By default, routers that support entropy labels are configured with the **load-balance-label-capability** statement at the **[edit forwarding-options]** hierarchy level to signal the labels on a per-LSP basis. If the peer router is not equipped to handle load-balancing labels, you can prevent the signaling of entropy label capability by configuring the **no-load-balance-label-capability** at the **[edit forwarding-options]** hierarchy level.

```
[edit forwarding-options]
user@PE# no-load-balance-label-capability
```

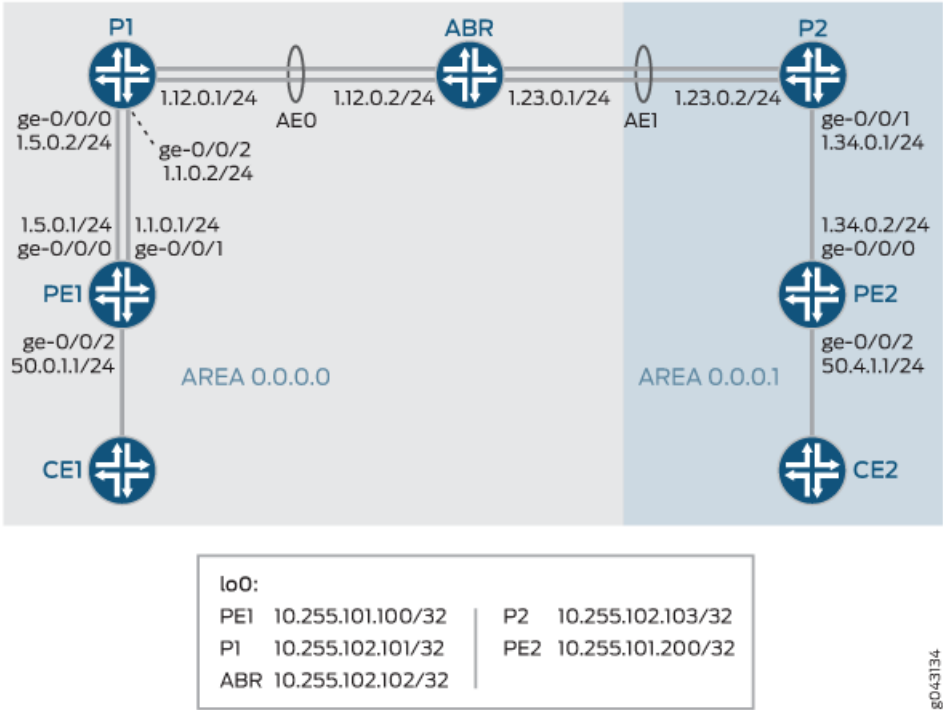
NOTE: You can explicitly disable advertising entropy label capability at egress for routes specified in the policy with the **no-entropy-label-capability** option at the **[edit policy-options policy-statement policy-name then]** hierarchy level.

```
[edit policy-options policy-statement policy-name then]
user@PE# no-entropy-label-capability
```

Topology

In [Figure 46 on page 577](#), Router PE1 is the ingress router and Router PE2 is the egress router. Routers P1 and P2 are the transit routers. Router ABR is the area bridge router between Area 0 and Area 1. LAG is configured on the provider routers for load balancing the traffic. Entropy label capability for BGP labeled unicast is enabled on the ingress Router PE1.

Figure 46: Configuring an Entropy Label for BGP Labeled Unicast



Configuration

IN THIS SECTION

- [Configuring Router PE1 | 583](#)
- [Configuring Router P1 | 587](#)
- [Configuring Router ABR | 589](#)
- [Results | 591](#)

CLI Quick Configuration

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

Router PE1


```

set interfaces ge-0/0/0 unit 0 family inet address 1.5.0.1/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6 address 2000::1:5:0:1/120
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 1.1.0.1/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6 address 2000::1:1:0:1/120
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 50.0.1.1/24
set interfaces ge-0/0/2 unit 0 family inet6 address 2000::1:34:0:2/120
set interfaces ge-0/0/3 vlan-tagging
set interfaces ge-0/0/3 unit 0 vlan-id 520
set interfaces ge-0/0/3 unit 0 family inet address 1.0.0.2/16
set interfaces lo0 unit 0 family inet address 10.255.101.100/32 primary
set routing-options router-id 10.255.101.100
set routing-options autonomous-system 1
set protocols rsvp interface all
set protocols mpls icmp-tunneling
set protocols mpls no-cspf
set protocols mpls label-switched-path r0-r2 to 10.255.102.102
set protocols mpls label-switched-path r0-r2 entropy-label
set protocols mpls interface all
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.101.100
set protocols bgp group ibgp family inet labeled-unicast entropy-label
set protocols bgp group ibgp neighbor 10.255.102.102 family inet labeled-unicast rib inet.3
set protocols bgp group ibgp neighbor 10.255.101.200 family inet-vpn unicast
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options prefix-list el-fec 10.255.101.200/32
set policy-options prefix-list el-fec-2 10.255.102.102/32
set policy-options policy-statement EL from prefix-list el-fec
set policy-options policy-statement EL then accept
set policy-options policy-statement EL-2 from prefix-list el-fec-2
set policy-options policy-statement EL-2 then accept
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf then accept
set policy-options policy-statement ospf-to-bgp from protocol ospf
set policy-options policy-statement ospf-to-bgp then accept
set policy-options policy-statement stat-to-bgp from protocol static

```



```

set policy-options policy-statement stat-to-bgp then accept
set policy-options community VPN members target:100:1
set routing-instances VPN-l3vpn instance-type vrf
set routing-instances VPN-l3vpn interface ge-0/0/2.0
set routing-instances VPN-l3vpn interface ge-0/0/3.0
set routing-instances VPN-l3vpn route-distinguisher 100.100.100.100:100
set routing-instances VPN-l3vpn vrf-target target:100:1
set routing-instances VPN-l3vpn routing-options static route 5.0.0.0/16 next-hop 1.0.0.1
set routing-instances VPN-l3vpn protocols ospf export bgp-to-ospf
set routing-instances VPN-l3vpn protocols ospf area 0.0.0.0 interface ge-0/0/2.0

```

Router P1

```

set interfaces ge-0/0/0 unit 0 family inet address 1.5.0.2/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6 address 2000::1:5:0:2/120
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/2 unit 0 family inet address 1.1.0.2/24
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family inet6 address 2000::1:1:0:2/120
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 gigether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 1.12.0.1/24
set interfaces ae0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.101/32 primary
set forwarding-options hash-key family mpls label-1
set forwarding-options hash-key family mpls label-2
set forwarding-options hash-key family mpls label-3
set forwarding-options enhanced-hash-key family mpls no-payload
set routing-options router-id 10.255.102.101
set routing-options autonomous-system 1
set routing-options forwarding-table export pplb
set protocols rsvp interface all
set protocols mpls icmp-tunneling
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all

```



```

set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set policy-options policy-statement pplb then load-balance per-packet

```

Router ABR

```

set interfaces ge-0/0/0 gigether-options 802.3ad ae0
set interfaces ge-0/0/1 gigether-options 802.3ad ae1
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 gigether-options 802.3ad ae1
set interfaces ae0 unit 0 family inet address 1.12.0.2/24
set interfaces ae0 unit 0 family mpls
set interfaces ae1 unit 0 family inet address 1.23.0.1/24
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.102/32 primary
set forwarding-options hash-key family mpls label-1
set forwarding-options hash-key family mpls label-2
set forwarding-options hash-key family mpls label-3
set forwarding-options enhanced-hash-key family mpls no-payload
set routing-options router-id 10.255.102.102
set routing-options autonomous-system 1
set routing-options forwarding-table export pplb
set protocols rsvp interface all
set protocols mpls icmp-tunneling
set protocols mpls label-switched-path r2-r0 to 10.255.101.100
set protocols mpls label-switched-path r2-r0 entropy-label
set protocols mpls label-switched-path r2-r4 to 10.255.101.200
set protocols mpls label-switched-path r2-r4 entropy-label
set protocols mpls interface all
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.102.102
set protocols bgp group ibgp family inet labeled-unicast rib inet.3
set protocols bgp group ibgp neighbor 10.255.101.100 export send-inet3-R4
set protocols bgp group ibgp neighbor 10.255.101.200 export send-inet3-R0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```



```

set protocols ospf area 0.0.0.1 interface ge-0/0/3.0
set protocols ospf area 0.0.0.1 interface ge-0/0/1.0
set protocols ospf area 0.0.0.1 interface ae1.0
set protocols ldp interface all
set policy-options policy-statement pplb then load-balance per-packet
set policy-options policy-statement send-inet3-R0 from route-filter 10.255.101.100/32 exact
set policy-options policy-statement send-inet3-R0 then accept
set policy-options policy-statement send-inet3-R4 from route-filter 10.255.101.200/32 exact
set policy-options policy-statement send-inet3-R4 then accept

```

Router P2

```

set chassis aggregated-devices ethernet device-count 3
set interfaces ge-0/0/0 gigether-options 802.3ad ae0
set interfaces ge-0/0/1 unit 0 family inet address 1.34.0.1/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6 address 2000::1:34:0:1/120
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 gigether-options 802.3ad ae0
set interfaces ae1 unit 0 family inet address 1.23.0.2/24
set interfaces ae1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.103/32 primary
set forwarding-options enhanced-hash-key family mpls no-payload
set routing-options router-id 10.255.102.103
set routing-options autonomous-system 1
set routing-options forwarding-table export pplb
set protocols rsvp interface all
set protocols mpls icmp-tunneling
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set protocols ospf area 0.0.0.1 interface fxp0.0 disable
set protocols ospf area 0.0.0.1 interface all
set policy-options policy-statement pplb then load-balance per-packet

```

Router PE2


```
set interfaces ge-0/0/0 unit 0 family inet address 1.34.0.2/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6 address 2000::1:34:0:2/120
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 vlan-tagging
set interfaces ge-0/0/1 unit 0 vlan-id 520
set interfaces ge-0/0/1 unit 0 family inet address 2.0.0.2/16
set interfaces ge-0/0/2 unit 0 family inet address 50.4.1.1/24
set interfaces ge-0/0/2 unit 0 family inet6 address 2000::1:34:0:2/120
set interfaces lo0 unit 0 family inet address 10.255.101.200/32 primary
set routing-options router-id 10.255.101.200
set routing-options autonomous-system 1
set protocols rsvp interface all
set protocols mpls icmp-tunneling
set protocols mpls no-cspf
set protocols mpls label-switched-path r4-r2 to 10.255.102.102
set protocols mpls label-switched-path r4-r2 entropy-label
set protocols mpls interface all
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.101.200
set protocols bgp group ibgp neighbor 10.255.102.102 family inet labeled-unicast rib inet.3
set protocols bgp group ibgp neighbor 10.255.101.100 family inet-vpn unicast
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.1 interface all
set protocols ospf area 0.0.0.1 interface fxp0.0 disable
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set policy-options prefix-list el-fec 10.255.101.100/32
set policy-options policy-statement EL term el from prefix-list el-fec
set policy-options policy-statement EL term el then accept
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf then accept
set policy-options policy-statement ospf-to-bgp from protocol ospf
set policy-options policy-statement ospf-to-bgp then accept
set policy-options policy-statement stat-to-bgp from protocol static
set policy-options policy-statement stat-to-bgp then accept
set policy-options community VPN members target:100:1
set routing-instances VPN-l3vpn instance-type vrf
set routing-instances VPN-l3vpn interface ge-0/0/1.0
set routing-instances VPN-l3vpn interface ge-0/0/2.0
set routing-instances VPN-l3vpn route-distinguisher 100.100.100.100:104
set routing-instances VPN-l3vpn vrf-target target:100:1
set routing-instances VPN-l3vpn routing-options static route 6.0.0.0/16 next-hop 2.0.0.1
```



```

set routing-instances VPN-l3vpn protocols ospf export bgp-to-ospf
set routing-instances VPN-l3vpn protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set routing-instances VPN-l3vpn protocols ospf area 0.0.0.0 interface ge-0/0/1.0

```

Configuring Router PE1

Step-by-Step Procedure

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

To configure Router PE1:

NOTE: Repeat this procedure for Router PE2 after modifying the appropriate interface names, addresses, and other parameters.

1. Configure the interfaces with IPv4 and IPv6 addresses.

```

[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 1.5.0.1/24
user@PE1# set ge-0/0/0 unit 0 family iso
user@PE1# set ge-0/0/0 unit 0 family inet6 address 2000::1:5:0:1/120
user@PE1# set ge-0/0/0 unit 0 family mpls

user@PE1# set ge-0/0/1 unit 0 family inet address 1.1.0.1/24
user@PE1# set ge-0/0/1 unit 0 family iso
user@PE1# set ge-0/0/1 unit 0 family inet6 address 2000::1:1:0:1/120
user@PE1# set ge-0/0/1 unit 0 family mpls

user@PE1# set ge-0/0/2 unit 0 family inet address 50.0.1.1/24
user@PE1# set ge-0/0/2 unit 0 family inet6 address 2000::1:34:0:2/120

user@PE1# set ge-0/0/3 vlan-tagging
user@PE1# set ge-0/0/3 unit 0 vlan-id 520
user@PE1# set ge-0/0/3 unit 0 family inet address 1.0.0.2/16

```

2. Configure the loopback interface.

```

[edit interfaces]

```



```
user@PE1# set lo0 unit 0 family inet address 10.255.101.100/32 primary
```

3. Set the router ID and the autonomous system number.

```
[edit routing-options]
user@PE1# set router-id 10.255.101.100
user@PE1# set autonomous-system 1
```

4. Configure RSVP protocol for all interfaces.

```
[edit protocols]
user@PE1# set protocols rsvp interface all
```

5. Enable MPLS on all the interfaces of Router PE1 and specify the LSP.

```
[edit protocols]
user@PE1# set mpls icmp-tunneling
user@PE1# set mpls no-cspf
user@PE1# set mpls label-switched-path r0-r2 to 10.255.102.102
user@PE1# set mpls label-switched-path r0-r2 entropy-label
user@PE1# set mpls interface all
```

6. Configure IBGP on the internal routers.

```
[edit protocols]
user@PE1# set bgp group ibgp type internal
user@PE1# set bgp group ibgp local-address 10.255.101.100
```

7. Enable entropy label capability for BGP labeled unicast for internal BGP group ibgp.

```
user@PE1# set bgp group ibgp family inet labeled-unicast entropy-label
user@PE1# set bgp group ibgp neighbor 10.255.102.102 family inet labeled-unicast rib inet.3
user@PE1# set bgp group ibgp neighbor 10.255.101.200 family inet-vpn unicast
```

8. Enable the OSPF protocol on all the interfaces of the area border router (ABR).

```
[edit protocols]
user@PE1# set ospf traffic-engineering
```



```

user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface lo0.0 passive

```

9. Define prefix lists to specify the routes with entropy label capability.

```

[edit policy-options ]
user@PE1# set policy-options prefix-list el-fec 10.255.101.200/32
user@PE1# set policy-options prefix-list el-fec-2 10.255.102.102/32

```

10. Define a policy EL to specify the routes with entropy label capability.

```

[edit policy-options ]
user@PE1# set policy-statement EL from prefix-list el-fec
user@PE1# set policy-statement EL then accept

```

11. Define another policy EL-2 to specify the routes with entropy label capability.

```

[edit policy-options ]
user@PE1# set policy-statement EL-2 from prefix-list el-fec-2
user@PE1# set policy-statement EL-2 then accept

```

12. Define a policy to export BGP routes to the OSPF routing table.

```

[edit policy-options ]
user@PE1# set policy-statement bgp-to-ospf from protocol bgp
user@PE1# set policy-statement bgp-to-ospf then accept

```

13. Define a policy to export OSPF routes to the BGP routing table.

```

[edit policy-options ]
user@PE1# set policy-statement ospf-to-bgp from protocol ospf
user@PE1# set policy-statement ospf-to-bgp then accept

```

14. Define a policy to export static routes to the BGP routing table.

```

[edit policy-options ]
user@PE1# set policy-statement stat-to-bgp from protocol static

```



```
user@PE1# set policy-statement stat-to-bgp then accept
```

15. Configure a VPN target for the VPN community.

```
[edit policy-options ]
user@PE1# set community VPN members target:100:1
```

16. Configure the Layer 3 VPN routing instance VPN-I3vpn.

```
[edit routing-instances]
user@PE1# set VPN-I3vpn instance-type vrf
```

17. Assign the interfaces for the VPN-I3vpn routing instance.

```
[edit routing-instances]
user@PE1# set VPN-I3vpn interface ge-0/0/2.0
user@PE1# set VPN-I3vpn interface ge-0/0/3.0
```

18. Configure the route distinguisher for the VPN-I3vpn routing instance.

```
[edit routing-instances]
user@PE1# set VPN-I3vpn route-distinguisher 100.100.100.100:100
```

19. Configure a VPN routing and forwarding (VRF) target for the VPN-I3vpn routing instance.

```
[edit routing-instances]
user@PE1# set VPN-I3vpn vrf-target target:100:1
```

20. Configure a static route to Device CE1 using the Layer 3 VPN protocol for the VPN-I3vpn routing instance.

```
[edit routing-instances]
user@PE1# set VPN-I3vpn routing-options static route 5.0.0.0/16 next-hop 1.0.0.1
```

21. Export the BGP routes to the OSPF routing table for the VPN-I3vpn routing instance.


```
[edit routing-instances]
user@PE1# set VPN-I3vpn protocols ospf export bgp-to-ospf
```

22. Assign the OSPF interface for the VPN-I3vpn routing instance.

```
[edit routing-instances]
user@PE1# set VPN-I3vpn protocols ospf area 0.0.0.0 interface ge-0/0/2.0
```

Configuring Router P1

Step-by-Step Procedure

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

To configure Router P1:

NOTE: Repeat this procedure for Router P2 after modifying the appropriate interface names, addresses, and other parameters.

1. Configure the interfaces with IPv4 and IPv6 addresses.

```
[edit interfaces]
user@P1# set ge-0/0/0 unit 0 family inet address 1.5.0.2/24
user@P1# set ge-0/0/0 unit 0 family iso
user@P1# set ge-0/0/0 unit 0 family inet6 address 2000::1:5:0:2/120
user@P1# set ge-0/0/0 unit 0 family mpls

user@P1# set ge-0/0/2 unit 0 family inet address 1.1.0.2/24
user@P1# set ge-0/0/2 unit 0 family iso
user@P1# set ge-0/0/2 unit 0 family inet6 address 2000::1:1:0:2/120
user@P1# set ge-0/0/2 unit 0 family mpls

user@P1# set ge-0/0/1 gigether-options 802.3ad ae0

user@P1# set ge-0/0/3 gigether-options 802.3ad ae0
```

2. Configure link aggregation on the interfaces.

```
user@P1# set ae0 unit 0 family inet address 1.12.0.1/24
```



```
user@P1# set ae0 unit 0 family mpls
```

3. Configure the loopback interface.

```
[edit interfaces]
user@P1# set lo0 unit 0 family inet address 10.255.102.101/32 primary
```

4. Configure MPLS labels that the router uses for hashing the packets to its destination for load balancing.

```
[edit forwarding-options]
user@P1# set hash-key family mpls label-1
user@P1# set hash-key family mpls label-2
user@P1# set hash-key family mpls label-3
user@P1# set enhanced-hash-key family mpls no-payload
```

5. Set the router ID and the autonomous system number.

```
[edit routing-options]
user@P1# set router-id 10.255.102.101
user@P1# set autonomous-system 1
```

6. Enable per packet load balancing.

```
[edit routing-options]
user@P1# set forwarding-table export pplb
```

7. Configure the RSVP protocol for all interfaces.

```
[edit protocols]
user@P1# set protocols rsvp interface all
```

8. Enable MPLS on all the interfaces of Router P1 and specify the LSP.

```
[edit protocols]
user@P1# set protocols mpls icmp-tunneling
user@P1# set protocols mpls interface all
```


9. Enable the OSPF protocol on all the interfaces of Router P1 excluding the management interface.

```
[edit protocols]
user@P1# set protocols ospf traffic-engineering
user@P1# set protocols ospf area 0.0.0.0 interface lo0.0 passive
user@P1# set protocols ospf area 0.0.0.0 interface fxp0.0 disable
user@P1# set protocols ospf area 0.0.0.0 interface all
user@P1# set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
user@P1# set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
```

10. Define a policy for per packet load balancing.

```
[edit policy-options]]
user@P1# set policy-statement pplb then load-balance per-packet
```

Configuring Router ABR

Step-by-Step Procedure

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

To configure Router ABR:

1. Configure the interfaces with IPv4 and IPv6 addresses.

```
[edit interfaces]
user@ABR# set ge-0/0/0 gigether-options 802.3ad ae0

user@ABR# set ge-0/0/1 gigether-options 802.3ad ae1

user@ABR# set ge-0/0/2 gigether-options 802.3ad ae0

user@ABR# set ge-0/0/3 gigether-options 802.3ad ae1
```

2. Configure the loopback interface.

```
[edit interfaces]
user@ABR# set lo0 unit 0 family inet address 10.255.102.102/32 primary
```

3. Configure link aggregation on the interfaces.


```
[edit interfaces]
user@ABR# set ae0 unit 0 family inet address 1.12.0.2/24
user@ABR# set ae0 unit 0 family mpls
user@ABR# set ae1 unit 0 family inet address 1.23.0.1/24
user@ABR# set ae1 unit 0 family mpls
```

4. Configure MPLS labels that the router uses for hashing the packets to its destination for load balancing.

```
[edit forwarding-options]
user@ABR# set hash-key family mpls label-1
user@ABR# set hash-key family mpls label-2
user@ABR# set hash-key family mpls label-3
user@ABR# set enhanced-hash-key family mpls no-payload
```

5. Set the router ID and the autonomous system number.

```
[edit routing-options]
user@ABR# set router-id 10.255.102.102
user@ABR# set autonomous-system 1
```

6. Enable per packet load balancing.

```
[edit routing-options]
user@ABR# set forwarding-table export pplb
```

7. Configure the RSVP protocol for all interfaces.

```
[edit protocols]
user@ABR# set protocols rsvp interface all
```

8. Enable MPLS on all the interfaces of Router P1 and specify the LSP.

```
[edit protocols]
user@ABR# set mpls icmp-tunneling
user@ABR# set mpls label-switched-path r2-r0 to 10.255.101.100
user@ABR# set mpls label-switched-path r2-r0 entropy-label
user@ABR# set mpls label-switched-path r2-r4 to 10.255.101.200
user@ABR# set mpls label-switched-path r2-r4 entropy-label
user@ABR# set mpls interface all
```


9. Configure IBGP on the internal routers.

```
[edit protocols ]
user@ABR# set bgp group ibgp type internal
user@ABR# set bgp group ibgp local-address 10.255.102.102
user@ABR# set bgp group ibgp family inet labeled-unicast rib inet.3
user@ABR# set bgp group ibgp neighbor 10.255.101.100 export send-inet3-R4
user@ABR# set bgp group ibgp neighbor 10.255.101.200 export send-inet3-R0
```

10. Enable the OSPF protocol on all the interfaces of ABR.

```
[edit protocols ]
user@ABR# set ospf traffic-engineering
user@ABR# set ospf area 0.0.0.0 interface lo0.0 passive
user@ABR# set ospf area 0.0.0.0 interface ge-0/0/2.0
user@ABR# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@ABR# set ospf area 0.0.0.0 interface ae0.0
user@ABR# set ospf area 0.0.0.0 interface fxp0.0 disable
user@ABR# set ospf area 0.0.0.1 interface ge-0/0/3.0
user@ABR# set ospf area 0.0.0.1 interface ge-0/0/1.0
user@ABR# set ospf area 0.0.0.1 interface ae1.0
```

11. Define a policy to specify the routes with entropy label capability.

```
[edit policy-options ]
user@ABR# set policy-statement pplb then load-balance per-packet
user@ABR# set policy-statement send-inet3-R0 from route-filter 10.255.101.100/32 exact
user@ABR# set policy-statement send-inet3-R0 then accept
user@ABR# set policy-statement send-inet3-R4 from route-filter 10.255.101.200/32 exact
user@ABR# set policy-statement send-inet3-R4 then accept
```

Results

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

```
[edit]
user@ABR# show interfaces
ge-0/0/0 {
  gigether-options {
```



```

        802.3ad ae0;
    }
}
ge-0/0/1 {
    gigeother-options {
        802.3ad ae1;
    }
}
ge-0/0/2 {
    gigeother-options {
        802.3ad ae0;
    }
}
ge-0/0/3 {
    gigeother-options {
        802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family inet {
            address 1.12.0.2/24;
        }
        family mpls;
    }
}
ae1 {
    unit 0 {
        family inet {
            address 1.23.0.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.102.102/32 {
                primary;
            }
        }
    }
}
}

```



```

[edit]
user@ABR# show protocols
rsvp {
    interface all;
}
mpls {
    icmp-tunneling;
    label-switched-path r2-r0 {
        to 10.255.101.100;
        entropy-label;
    }
    label-switched-path r2-r4 {
        to 10.255.101.200;
        entropy-label;
    }
    interface all;
}
bgp {
    group ibgp {
        type internal;
        local-address 10.255.102.102;
        family inet {
            labeled-unicast {
                rib {
                    inet.3;
                }
            }
        }
        neighbor 10.255.101.100 {
            export send-inet3-R4;
        }
        neighbor 10.255.101.200 {
            export send-inet3-R0;
        }
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface ge-0/0/2.0;
        interface ge-0/0/0.0;
    }
}

```



```

    interface ae0.0;
    interface fxp0.0 {
        disable;
    }
}
area 0.0.0.1 {
    interface ge-0/0/3.0;
    interface ge-0/0/1.0;
    interface ae1.0;
}
}

```

```

[edit]
user@ABR# show routing-options
router-id 10.255.102.102;
autonomous-system 1;
forwarding-table {
    export pplb;
}

```

```

[edit]
user@ABR# show forwarding-options
hash-key {
    family mpls {
        label-1;
        label-2;
        label-3;
    }
}
enhanced-hash-key {
    family mpls {
        no-payload;
    }
}

```

```

[edit]
user@ABR# show policy-options
policy-statement pplb {
    then {
        load-balance per-packet;
    }
}

```



```

policy-statement send-inet3-R0 {
  from {
    route-filter 10.255.101.100/32 exact;
  }
  then accept;
}
policy-statement send-inet3-R4 {
  from {
    route-filter 10.255.101.200/32 exact;
  }
  then accept;
}

```

Verification

IN THIS SECTION

- [Verifying That the Entropy Label Capability Is Being Advertised from Router PE2 | 595](#)
- [Verifying That Router ABR Receives the Entropy Label Advertisement | 596](#)
- [Verifying That the Entropy Label Flag Is Set | 598](#)

Confirm that the configuration is working properly.

Verifying That the Entropy Label Capability Is Being Advertised from Router PE2

Purpose

Verify that the entropy label capability path attribute is being advertised from the upstream Router PE2 at egress.

Action

From operational mode, run the **show route 10.255.101.200 advertising-protocol bgp 10.255.102.102** command on Router PE2.

```
user@PE2> show route 10.255.101.200 advertising-protocol bgp 10.255.102.102
```

```

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.101.200/32 (1 entry, 1 announced)
  BGP group ibgp type Internal

```



```

Route Label: 299920
NextHop: Self
Flags: NextHop Change
MED: 2
Localpref: 4294967294
AS path: [1] I
Entropy label capable

```

Meaning

The output shows that the host PE2 with the IP address of 10.255.101.200 has the entropy label capability. The host is advertising the entropy label capability to its BGP neighbors.

Verifying That Router ABR Receives the Entropy Label Advertisement

Purpose

Verify that Router ABR receives the entropy label advertisement at ingress from Router PE2.

Action

From operational mode, run the **show route 10.255.101.200 receiving-protocol bgp 10.255.101.200** command on Router ABR.

```
user@ABR> show route 10.255.101.200 receiving-protocol bgp 10.255.101.200
```

```

inet.0: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.101.100/32 (1 entry, 1 announced)
  Accepted
  Route Label: 299920
  NextHop: 10.255.102.102
  MED: 2
  Localpref: 4294967294
  AS path: I
  Entropy label capable

VPN-13vpn.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

iso.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.13vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```



```

inet6.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)

VPN-l3vpn.inet6.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

user@PE1> show route protocol bgp detail

inet.0: 64 destinations, 64 routes (64 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.255.101.200/32 (1 entry, 1 announced)
    *BGP      Preference: 170/1
              Next hop type: Indirect, Next hop index: 0
              Address: 0xa533c10
              Next-hop reference count: 2
              Source: 10.255.102.102
              Next hop type: Router, Next hop index: 0
              Next hop: 1.1.0.2 via ge-0/0/1.0, selected
              Label-switched-path r0-r2
              Label operation: Push 299904, Push 300096(top)
              Label TTL action: prop-ttl, prop-ttl(top)
              Load balance label: Label 299904: Entropy label; Label 300096: None;
              Label element ptr: 0xa5335a0
              Label parent element ptr: 0xa5338a0
              Label element references: 2
              Label element child references: 1
              Label element lsp id: 0
              Session Id: 0x0
              Protocol next hop: 10.255.102.102
              Label operation: Push 299904
              Label TTL action: prop-ttl
              Load balance label: Label 299904: Entropy label;
              Indirect next hop: 0xaa18540 - INH Session ID: 0x0
              State: <Active Int Ext>
              Local AS:      1 Peer AS:      1
              Age: 12:39      Metric: 2      Metric2: 2
              Validation State: unverified
              Task: BGP_1.10.255.102.102
              Announcement bits (2): 0-Resolve tree 1 3-Resolve_IGP_FRR task

              AS path: I
              Accepted
              Route Label: 299904
              Localpref: 4294967294

```



```
Router ID: 10.255.102.102
VPN-l3vpn.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
```

Meaning

Router ABR receives the entropy label capability advertisement from its BGP neighbor PE2.

Verifying That the Entropy Label Flag Is Set

Purpose

Verify that the entropy label flag is set for the label elements at the ingress.

Action

From operational mode, run the **show route protocol bgp detail** command on Router PE1.

```
user@PE1> show route protocol bgp detail
```

```
inet.0: 64 destinations, 64 routes (64 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.255.101.200/32 (1 entry, 1 announced)
    *BGP      Preference: 170/1
              Next hop type: Indirect, Next hop index: 0
              Address: 0xa533c10
              Next-hop reference count: 2
              Source: 10.255.102.102
              Next hop type: Router, Next hop index: 0
              Next hop: 1.1.0.2 via ge-0/0/1.0, selected
              Label-switched-path r0-r2
              Label operation: Push 299904, Push 300096(top)
              Label TTL action: prop-ttl, prop-ttl(top)
              Load balance label: Label 299904: Entropy label; Label 300096: None;
              Label element ptr: 0xa5335a0
              Label parent element ptr: 0xa5338a0
              Label element references: 2
              Label element child references: 1
              Label element lsp id: 0
              Session Id: 0x0
              Protocol next hop: 10.255.102.102
              Label operation: Push 299904
              Label TTL action: prop-ttl
              Load balance label: Label 299904: Entropy label;
              Indirect next hop: 0xaa18540 - INH Session ID: 0x0
              State:      <Active Int Ext>
```



```

Local AS:      1 Peer AS:      1
Age: 12:39      Metric: 2      Metric2: 2
Validation State: unverified
Task: BGP_1.10.255.102.102
Announcement bits (2): 0-Resolve tree 1 3-Resolve_IGP_FRR task

AS path: I
Accepted
Route Label: 299904
Localpref: 4294967294
Router ID: 10.255.102.102
VPN-l3vpn.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

```

Meaning

An entropy label is enabled on Router PE1. The output shows that the entropy label is being used for the BGP labeled unicast to achieve end-to-end load balancing.

RELATED DOCUMENTATION

[entropy-label](#) | **2603**


Configuring an Entropy Label for a BGP Labeled Unicast LSP

Understanding Entropy Label for BGP Labeled Unicast LSP

Configuring Ultimate-Hop Popping for LSPs


By default, RSVP-signaled LSPs use penultimate-hop popping (PHP). [Figure 47 on page 600](#) illustrates a penultimate-hop popping LSP between Router PE1 and Router PE2. Router CE1 forwards a packet to its next hop (Router PE1), which is also the LSP ingress. Router PE1 pushes label 1 on the packet and forwards the labeled packet to Router P1. Router P1 completes the standard MPLS label swapping operation, swapping label 1 for label 2, and forwards the packet to Router P2. Since Router P2 is the penultimate-hop router for the LSP to Router PE2, it first pops the label and then forwards the packet to Router PE2. When Router PE2 receives it, the packet can have a service label, an explicit-null label, or just be a plain IP or VPLS packet. Router PE2 forwards the unlabeled packet to Router CE2.

Figure 47: Penultimate-Hop Popping for an LSP



You can also configure ultimate-hop popping (UHP) (as shown in [Figure 48 on page 600](#)) for RSVP-signaled LSPs. Some network applications can require that packets arrive at the egress router (Router PE2) with a non-null outer label. For an ultimate-hop popping LSP, the penultimate router (Router P2 in [Figure 48 on page 600](#)) performs the standard MPLS label swapping operation (in this example, label 2 for label 3) before forwarding the packet to egress Router PE2. Router PE2 pops the outer label and performs a second lookup of the packet address to determine the end destination. It then forwards the packet to the appropriate destination (either Router CE2 or Router CE4).

Figure 48: Ultimate-Hop Popping for an LSP



The following network applications require that you configure UHP LSPs:

- MPLS-TP for performance monitoring and in-band OAM
- Edge protection virtual circuits

The following features do not support the UHP behavior:

- LDP-signaled LSPs
- Static LSPs
- Point-to-multipoint LSPs
- CCC
- **traceroute** command

For more information about UHP behavior, see Internet draft [draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt](#), *Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs*.

For point-to-point RSVP-signaled LSPs, UHP behavior is signaled from the LSP ingress. Based on the ingress router configuration, RSVP can signal the UHP LSP with the non-PHP flag set. RSVP PATH messages carry the two flags in the LSP-ATTRIBUTES object. When the egress router receives the PATH message, it assigns a non-null label to the LSP. RSVP also creates and installs two routes in the mpls.0 routing table. S refers to the S bit of the MPLS label, which indicates whether or not the bottom of the label stack has been reached.

- Route S=0—Indicates that there are more labels in the stack. The next hop for this route points to the mpls.0 routing table, triggering a chained MPLS label lookup to discover the remaining MPLS labels in the stack.
- Route S=1—Indicates that there are no more labels. The next hop points to the inet.0 routing table if the platform supports chained and multi-family lookup. Alternatively, the label route can point to a VT interface to initiate IP forwarding.

If you enable UHP LSPs, MPLS applications such as Layer 3 VPNs, VPLS, Layer 2 VPNs, and Layer 2 circuits can use the UHP LSPs. The following explains how UHP LSPs affect the different types of MPLS applications:

- Layer 2 VPNs and Layer 2 circuits—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label (S=1) is the VC label. A lookup based on the transport label results in a table handle for the mpls.0 routing table. There is an additional route in the mpls.0 routing table corresponding to the inner label. A lookup based on the inner label results in the CE router next hop.
- Layer 3 VPN—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label is the VPN label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. There are two cases in this scenario. By default, Layer 3 VPNs advertise the per-next hop label. A lookup based on the inner label results in the next hop toward the CE router. However, if you have configured the **vrf-table-label** statement for the Layer 3 VPN routing instance, the inner LSI label points to the VRF routing table. An IP lookup is also completed for the VRF routing table.

NOTE: UHP for Layer 3 VPNs configured with the **vrf-table-label** statement is supported on MX Series 5G Universal Routing Platforms only.

- VPLS—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0) and the inner label is the VPLS label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. A lookup based on the inner label in mpls.0 routing table results in the LSI tunnel interface of the VPLS routing instance if tunnel-services is not configured (or a VT interface not available). MX 3D Series routers support chained lookup and multi-family lookup.

NOTE: UHP for VPLS configured with the **no-tunnel-service** statement is supported on MX 3D Series routers only.

- **IPv4 over MPLS**—A packet arrives at the PE router (egress of the UHP LSP) with one label (S=1). A lookup based on this label returns a VT tunnel interface. Another IP lookup is completed on the VT interface to determine where to forward the packet. If the routing platform supports multi-family and chained lookups (for example, MX 3D routers and PTX Series Packet Transport Routers), lookup based on label route (S=1) points to the inet.0 routing table.
- **IPv6 over MPLS**—For IPv6 tunneling over MPLS, PE routers advertise IPv6 routes to each other with a label value of 2. This is the explicit null label for IPv6. As a result, the forwarding next hops for IPv6 routes that are learned from remote PE routers normally push two labels. The inner label is 2 (it could be different if the advertising PE router is from another vendor), and the router label is the LSP label. Packets arrive at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0), and the inner label is the IPv6 explicit-null label (label 2). Lookup based on the inner label in the mpls.0 routing table redirects back to the mpls.0 routing table. On MX 3D Series routers, the inner label (label 2) is stripped off and an IPv6 lookup is done using the inet6.0 routing table.
- **Enabling both PHP and UHP LSPs**—You can configure both PHP and UHP LSPs over the same network paths. You can separate PHP and UHP traffic by selecting forwarding LSP next hops using a regular expression with the **install-nexthop** statement. You can also separate traffic by simply naming the LSPs appropriately.

The following statements enable ultimate-hop popping for an LSP. You can enable this feature on a specific LSP or for all of the ingress LSPs configured on the router. Configure these statements on the router at the LSP ingress.

1. To enable ultimate-hop popping, include the **ultimate-hop-popping** statement:

```
ultimate-hop-popping;
```

Include this statement at the **[edit protocols mpls label-switched-path label-switched-path-name]** hierarchy level to enable ultimate-hop popping on a specific LSP. Include this statement at the **[edit protocols mpls]** hierarchy level to enable ultimate-hop popping on all of the ingress LSPs configured on the router. You can also configure the **ultimate-hop-popping** statement under the equivalent **[edit logical-routers]** hierarchy levels.

NOTE: When you enable ultimate-hop popping, RSVP attempts to resignal existing LSPs as ultimate-hop popping LSPs in a make-before-break fashion. If an egress router does not support ultimate-hop popping, the existing LSP is torn down (RSVP sends a PathTear message along an LSP's path, removing the path state and dependent reservation state and releasing the associated networking resources).

If you disable ultimate-hop popping, RSVP resignals existing LSPs as penultimate-hop popping LSPs in a make-before-break fashion.

2. If you want to enable both ultimate-hop-popping and chained next hops on MX 3D Series routers only, you also need to configure the **enhanced-ip** option for the **network-services** statement:

```
network-services enhanced-ip;
```

You configure this statement at the **[edit chassis]** hierarchy level. Once you have configured the **network-services** statement, you need to reboot the router to enable UHP behavior.

RELATED DOCUMENTATION

[MPLS Label Allocation | 420](#)

[Configuring Corouted Bidirectional LSPs | 570](#)

[network-services](#)

[ultimate-hop-popping | 2824](#)

Configuring Static LSPs

IN THIS SECTION

- [Configuring the Ingress Router for Static LSPs | 604](#)
- [Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs | 607](#)
- [Configuring a Bypass LSP for the Static LSP | 610](#)
- [Configuring the Protection Revert Timer for Static LSPs | 610](#)
- [Configuring Static Unicast Routes for Point-to-Multipoint LSPs | 611](#)

To configure static LSPs, configure the ingress router and each router along the path up to and including the egress router.

To configure static MPLS, perform the following tasks:

Configuring the Ingress Router for Static LSPs

The ingress router checks the IP address in the incoming packet's destination address field and, if it finds a match in the routing table, applies the label associated with that address to the packets. The label has forwarding information associated with it, including the address of the next-hop router, and the route preference and CoS values.

To configure static LSPs on the ingress router, include the **ingress** statement:

```
ingress {
  bandwidth bps;
  class-of-service cos-value;
  description string;
  install {
    destination-prefix <active>;
  }
  link-protection bypass-name name;
  metric metric;
  next-hop (address | interface-name | address/interface-name);
  no-install-to-address;
  node-protection bypass-name name next-next-label label;
  policing {
    filter filter-name;
    no-auto-policing;
  }
  preference preference;
  push out-label;
  to address;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls static-label-switched-path *static-lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *static-lsp-name*]

When you configure a static LSP on the ingress router, the **next-hop**, **push**, and **to** statements are required; the other statements are optional.

The configuration for a static LSP on the ingress router requires you to configure the following parts:

- Criteria for analyzing an incoming packet:
 - The **install** statement creates an LSP that handles IPv4 packets. All static MPLS routes created using the **install** statement are installed in inet.3 routing table, and the creating protocol is identified as static. This process is no different from creating static IPv4 routes at the **[edit routing-options static]** hierarchy level.
 - In the **to** statement, you configure the IP destination address to check when incoming packets are analyzed. If the address matches, the specified outgoing label (**push out-label**) is assigned to the packet, and the packet enters an LSP. Manually assigned outgoing labels can have values from 0 through 1,048,575. Each prefix that you specify is installed as a static route in the routing table.
- The **next-hop** statement, which supplies the IP address of the next hop to the destination. You can specify this as the IP address of the next hop, the interface name (for point-to-point interfaces only), or as **address/interface-name** to specify an IP address on an operational interface. When the next hop is on a directly attached interface, the route is installed in the routing table. You cannot configure a LAN or nonbroadcast multiaccess (NBMA) interface as a next-hop interface.
- Properties to apply to the LSP (all are optional):
 - Bandwidth reserved for this LSP (**bandwidth bps**)
 - Link protection and node protection to apply to the LSP (**bypass bypass-name, link-protection bypass-name name, node-protection bypass-name next-next-label label**)
 - Metric value to apply to the LSP (**metric**)
 - Class-of-service value to apply to the LSP (**class-of-service**)
 - Preference value to apply to the LSP (**preference**)
 - Traffic policing to apply to the LSP (**policing**)
 - Text description to apply to the LSP (**description**)
 - Install or no-install policy (**install** or **no-install-to-address**)

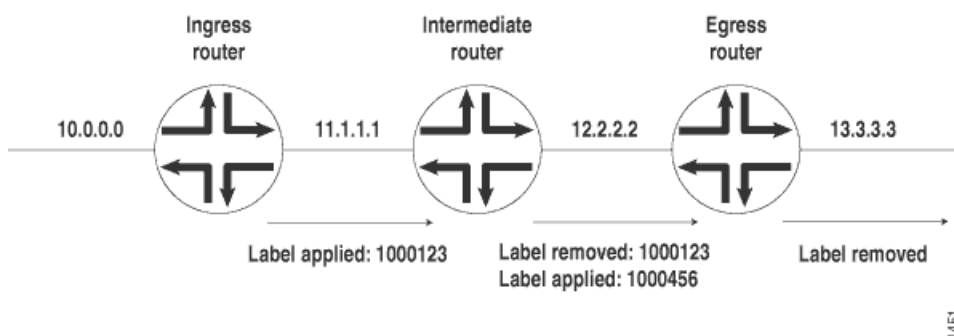
To determine whether a static ingress route is installed, use the command **show route table inet.3 protocol static**. Sample output follows. The **push** keyword denotes that a label is to be added in front of an IP packet.

```
10.0.0.0      *[Static/5] 00:01:48
> to 11.1.1.1 via so-0/0/0, push 1000123
```

Example: Configuring the Ingress Router

Configure the ingress router for a static LSP that consists of three routers (see [Figure 49 on page 606](#)).

Figure 49: Static MPLS Configuration



For packets addressed to **10.0.0.0**, assign label **1000123** and transmit them to the next-hop router at **11.1.1.1**:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  mpls {
    static-label-switched-path path1 {
      ingress {
        next-hop 11.1.1.1;
        to 10.0.0.0;
        push 1000123;
      }
    }
    interface so-0/0/0.0;
  }
}
routing-options {
  static {
    route 10.0.0.0/8 {
      static-lsp-next-hop path1;
    }
  }
}
```

To determine whether the static ingress route is installed, use the following command:

```
user@host> show route table inet.0 protocol static
```


Sample output follows. The **push 1000123** keyword identifies the route.

```
10.0.0.0/8          *[Static/5] 00:01:48
> to 11.1.1.1 via so-0/0/0.0, push 1000123
```

Configuring the Intermediate (Transit) and Egress Routers for Static LSPs

Intermediate (transit) and egress routers perform similar functions—they modify the label that has been applied to a packet. An intermediate router can change the label. An egress router removes the label (if the packet still contains a label) and continues forwarding the packet to its destination.

To configure static LSPs on intermediate and egress routers, include the **transit** statement:

```
static-label-switched-path lsp-name {
  transit incoming-label {
    bandwidth bps;
    description string;
    link-protection bypass-name name;
    next-hop (address | interface-name | address/interface-name);
    node-protection bypass-name name next-next-label label;
    pop;
    swap out-label;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls static-label-switched-path *static-lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *static-lsp-name*]

For the **transit** statement configuration, the **next-hop** and **pop** | **swap** statements are required. The remaining statements are optional.

Each statement within the **transit** statement consists of the following parts:

- Packet label (specified in the **transit** statement)
- The **next-hop** statement, which supplies the IP address of the next hop to the destination. The address is specified as the IP address of the next hop, or the interface name (for point-to-point interfaces only), or **address** and **interface-name** to specify an IP address on an operational interface. When the specified next hop is on a directly attached interface, this route is installed in the routing table. You cannot configure a LAN or NBMA interface as a next-hop interface.

- Operation to perform on the labeled packet:
 - For egress routers, you generally just remove the packet's label altogether (**pop**) and continue forwarding the packet to the next hop. However, if the previous router removed the label, the egress router examines the packet's IP header and forwards the packet toward its IP destination.
 - For intermediate (transit) routers only, exchange the label for another label (**swap out-label**). Manually assigned incoming labels can have values from 1,000,000 through 1,048,575. Manually assigned outgoing labels can have values from 0 through 1,048,575.
- Label properties to apply to the packet (all are optional):
 - Bandwidth reserved for this route (**bandwidth bps**).
 - Link-protection and node-protection to apply to the LSP (**bypass bypass-name, link-protection bypass-name name, node-protection bypass-name next-next-label label**).
 - Text description to apply to the LSP (specified in the **description** statement).

The static routes are installed in the default MPLS routing table, `mpls.0`, and the creating protocol is identified as static. To verify that a static route is properly installed, use the command **show route table mpls.0 protocol static**. Sample output follows:

```
1000123      *[Static/5] 00:00:38
> to 12.2.2.2 via so-5/0/0.0, swap 1000456
```

You can configure a revert timer for a static LSP transiting an intermediate router. After traffic has been switched to a bypass static LSP, it is typically switched back to the primary static LSP when it comes back up. There is a configurable delay in the time (called the revert timer) between when the primary static LSP comes up and when traffic is reverted back to it from the bypass static LSP. This delay is needed because when the primary LSP comes back up, it is not certain whether all of the interfaces on the downstream node of the primary path have come up yet. You can display the revert timer value for an interface using the **show mpls interface detail** command. For more information, see [“Configuring the Revert Timer for LSPs” on page 567](#).

Example: Configuring an Intermediate Router

For packets labeled **1000123** arriving on interface **so-0/0/0**, assign the label **1000456**, and transmit them to the next-hop router at **12.2.2.2**:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
```



```

    }
}
protocols {
  mpls {
    static-label-switched-path path1 {
      transit 1000123 {
        next-hop 12.2.2.2;
        swap 1000456;
      }
    }
  }
  interface so-0/0/0.0;
}
}

```

To determine whether the static intermediate route is installed, use the following command:

```
user@host> show route table mpls.0 protocol static
```

Sample output follows. The **swap 1000456** keyword identifies the route.

```

1000123      *[Static/5] 00:01:48
> to 12.2.2.2 via so-0/0/0, swap 1000456

```

Example: Configuring an Egress Router

For packets labeled **1000456** arriving on interface **so-0/0/0**, remove the label and transmit the packets to the next-hop router at **13.3.3.3**:

```

[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  mpls {
    static-label-switched-path path1 {
      transit 1000456 {
        next-hop 13.3.3.3;
        pop;
      }
    }
  }
}

```



```

    }
  }
  interface so-0/0/0.0;
}
}

```

To determine whether the static egress route is installed, use the following command:

```
user@host> show route table mpls.0 protocol static
```

Sample output follows. The **pop** keyword identifies the egress route.

```

1000456          *[Static/5] 00:01:48
> to 13.3.3.3 via so-0/0/0, pop

```

Configuring a Bypass LSP for the Static LSP

To enable a bypass LSP for the static LSP, configure the **bypass** statement:

```

bypass bypass-name {
  bandwidth bps;
  description string;
  next-hop (address | interface-name | address/interface-name);
  push out-label;
  to address;
}

```

Configuring the Protection Revert Timer for Static LSPs

For static LSPs configured with a bypass static LSP, it is possible to configure the protection revert timer. If a static LSP goes down and traffic is switched to the bypass LSP, the protection revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert back to the original static LSP.

The range of values you can configure for the protection revert timer is 0 through 65,535 seconds. The default value is 5 seconds.

If you configure a value of 0 seconds, the traffic on the LSP, once switched from the original static LSP to the bypass static LSP, remains on the bypass LSP permanently (until the network operator intervenes or until the bypass LSP goes down).

You can configure the protection revert timer for all LSPs on the router at the **[edit protocols mpls]** hierarchy level or for a specific LSP at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level.

To configure the protection revert timer for static LSPs include the **protection-revert-time** statement:

```
protection-revert-time seconds;
```

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

Configuring Static Unicast Routes for Point-to-Multipoint LSPs

You can configure a static unicast IP route with a point-to-multipoint LSP as the next hop. For more information about point-to-multipoint LSPs, see [“Point-to-Multipoint LSPs Overview” on page 691](#), [“Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs” on page 720](#), and [“Configuring CCC Switching for Point-to-Multipoint LSPs” on page 1681](#).

To configure a static unicast route for a point-to-multipoint LSP, complete the following steps:

1. On the ingress PE router, configure a static IP unicast route with the point-to-multipoint LSP name as the next hop by including the **p2mp-lsp-next-hop** statement:

```
p2mp-lsp-next-hop point-to-multipoint-lsp-next-hop;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options static route *route-name*]**
 - **[edit logical-systems *logical-system-name* routing-options static route *route-name*]**
2. On the egress PE router, configure a static IP unicast route with the same destination address configured in Step 1 (the address configured at the **[edit routing-options static route]** hierarchy level) by including the **next-hop** statement:

```
next-hop address;
```

You can include this statement at the following hierarchy levels:

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

NOTE: CCC and static routes cannot use the same point-to-multipoint LSP.

For more information on static routes, see the *Junos OS Routing Protocols Library*.

The following **show route** command output displays a unicast static route pointing to a point-to-multipoint LSP on the ingress PE router where the LSP has two branch next hops:

```
user@host> show route 5.5.5.5 detail
```

```
inet.0: 29 destinations, 30 routes (28 active, 0 holddown, 1 hidden)
5.5.5.5/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next hop type: Flood
    Next hop: via so-0/3/2.0 weight 1
    Label operation: Push 100000
    Next hop: via tl-0/1/1.0 weight 1
    Label operation: Push 100064
    State: <Active Int Ext>
    Local AS: 10458
    Age: 2:41:15
    Task: RT
    Announcement bits (2): 0-KRT 3-BGP.0.0.0.0+179
    AS path: I
```

Configuring Static Label Switched Paths for MPLS (CLI Procedure)

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, liveness detection, or statistics reporting.

To configure static LSPs, configure the ingress switch and each provider switch along the path up to and including the egress switch.

For the ingress 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. Optionally, you can apply preference, class-of-service (CoS) values, node protection, and link protection to the packets.

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. Optionally, you can apply node protection and link protection to the packets.

For the egress switch, you generally just remove the label and continue forwarding 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 an LSP, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 95.](#)
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99.](#)

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 | 613](#)
2. [Configuring the Provider and the Egress PE Switch | 614](#)

Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the core interfaces:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
```

2. Configure the name and the traffic rate associated with the LSP:

```
[edit]
user@switch# set protocols mpls static-label-switched-path lsp-name ingress bandwidth rate
```

3. Configure the next hop switch for the LSP:

```
[edit]
user@switch# set protocols mpls static-label-switched-path lsp-name ingress next-hop
address-of-next-hop
```

4. Enable link protection on the specified static LSP:

```
[edit]
user@switch# set protocols mpls static-label-switched-path lsp-name ingress link-protection
bypass-name name
```

5. Specify the address of the egress switch for the LSP:

```
[edit]
user@switch# set protocols mpls static-label-switched-path path1 ingress to address-of-egress-switch
```


6. Configure the new label that you want to add to the top of the label stack:

```
[edit]
user@switch# set protocols mpls static-label-switched-path path1 ingress push out-label
```

7. Optionally, configure the next hop address and the egress router address that you want to bypass, for the static LSP:

```
[edit]
user@switch# set protocols mpls static-label-switched-path lsp-name by bypass next-hop
address-of-next-hop
user@switch# set protocols mpls static-label-switched-path lsp-name by bypass to
address-of-the-egress-switch
user@switch# set protocols mpls static-label-switched-path lsp-name bypass push out-label
```

Configuring the Provider and the Egress PE Switch

To configure a static LSP for MPLS on the provider and egress provider edge switch:

1. Configure a transit static LSP:

```
[edit]
user@switch# set protocols mpls static-label-switched-path path1 transit incoming-label
```

2. Configure the next hop switch for the LSP:

```
[edit]
user@switch# set protocols mpls 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]
user@switch# set protocols mpls static-label-switched-path lsp-name transit incoming-label swap
out-label
```

4. Only for the egress provider edge 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]

```
user@switch# set protocols mpls static-label-switched-path lsp-name transit incoming-label pop
```

RELATED DOCUMENTATION

| [Example: Configuring MPLS on EX8200 and EX4500 Switches](#) | 59

Enabling Distributed CSPF for Segment Routing LSPs

IN THIS SECTION

- [Distributed CSPF Computation Constraints](#) | 616
- [Distributed CSPF Computation Algorithm](#) | 616
- [Distributed CSPF Computation Database](#) | 617
- [Configuring Distributed CSPF Computation Constraints](#) | 617
- [Distributed CSPF Computation](#) | 619
- [Interaction Between Distributed CSPF Computation and SRTE Features](#) | 619
- [Distributed CSPF Computation Sample Configurations](#) | 620

Prior to Junos OS Release 19.2R1S1, for traffic engineering of segment routing paths, you could either explicitly configure static paths, or use computed paths from an external controller. With the distributed Constrained Shortest Path First (CSPF) for segment routing LSP feature, you can compute a segment routing LSP locally on the ingress device according to the constraints you have configured. With this feature, the LSPs are optimized based on the configured constraints and metric type (traffic-engineering or IGP). The LSPs are computed to utilize the available ECMP paths to the destination with segment routing label stack compression enabled or disabled.

Distributed CSPF Computation Constraints

Segment routing LSP paths are computed when all the configured constraints are met.

The distributed CSPF computation feature supports the following subset of constraints specified in the Internet draft, draft-ietf-spring-segment-routing-policy-03.txt, *Segment Routing Policy for Traffic Engineering*:

- Inclusion and exclusion of administrative groups.
- Inclusion of loose or strict hop IP addresses.

NOTE: You can specify only router IDs in the loose or strict hop constraints. Labels and other IP addresses cannot be specified as loose or strict hop constraints in Junos OS Release 19.2R1-S1.

- Maximum number of segment IDs (SIDs) in the segment list.
- Maximum number of segment lists per candidate segment routing path.

The distributed CSPF computation feature for segment routing LSPs does not support the following types of constraints and deployment scenarios:

- IPV6 addresses.
- Inter domain segment routing traffic engineering (SRTE) LSPs.
- Unnumbered interfaces.
- Multiple protocols routing protocols such as, OSPF, ISIS, and BGP-LS, enabled at the same time.
- Computation with prefixes or anycast addresses as destinations.
- Including and excluding interface IP addresses as constraints.

Distributed CSPF Computation Algorithm

IN THIS SECTION

- [Label Stack Compression Enabled | 617](#)
- [Label Stack Compression Disabled | 617](#)

The distributed CSPF computation feature for segment routing LSPs uses the label stack compression algorithm with CSPF.

Label Stack Compression Enabled

A compressed label stack represents a set of paths from a source to a destination. It generally consists of node SIDs and adjacency SIDs. When label stack compression is enabled, the result of the computation is a set of paths that maximize ECMP to the destination, with minimum number of SIDs in the stack, while conforming to constraints.

Label Stack Compression Disabled

The multipath CSPF computation with label stack compression disabled finds up to N segment lists to destination, where:

- The cost of all segment lists is equal to and the same as the shortest traffic-engineering metric to reach the destination.
- Each segment list is comprised of adjacency SIDs.
- The value of N is the maximum number of segment lists allowed for the candidate path by configuration.
- No two segment lists are identical.
- Each segment list satisfies all the configured constraints.

Distributed CSPF Computation Database

The database used for SRTE computation has all links, nodes, prefixes and their characteristics irrespective of whether traffic-engineering is enabled in those advertising nodes. In other words, it is the union of the traffic-engineering database (TED) and the IGP link state database of all domains that the computing node has learnt from.

Configuring Distributed CSPF Computation Constraints

You can use a compute profile to logically group the computation constraints. These compute profiles are referenced by the segment routing paths for computing the primary and secondary segment routing LSPs.

To configure a compute profile, include the [compute-profile](#) statement at the **[edit protocols source-packet-routing]** hierarchy level.

The configuration for the supported computation constraints include:

- **Administrative groups**

You can configure [admin-groups](#) under the **[edit protocols mpls]** hierarchy level. Junos OS applies the administrative group configuration to the segment routing traffic-engineering (SRTE) interfaces.

To configure the computation constraints you can specify three categories for a set of administrative groups. The computation constraint configuration can be common to all candidate segment routing paths, or it can be under individual candidate paths.

- **include-any**—Specifies that any link with at least one of the configured administrative groups in the list is acceptable for the path to traverse.
- **include-all**—Specifies that any link with all of the configured administrative groups in the list is acceptable for the path to traverse.
- **exclude**—Specifies that any link which does not have any of the configured administrative groups in the list is acceptable for the path to traverse.

- **Explicit path**

You can specify a series of router IDs in the compute profile as a constraint for computing the SRTE candidate paths. Each hop has to be an IPv4 address and can be of type strict or loose. If the type of a hop is not configured, strict is used. You must include the **compute** option under the **segment-list** statement when specifying the explicit path constraint.

- **Maximum number of segment lists (ECMP paths)**

You can associate a candidate path with a number of dynamic segment-lists. The paths are ECMP paths, where each segment-list translates into a next hop gateway with active weight. These paths are a result of path computation with or without compression.

You can configure this attribute using the **maximum-computed-segment-lists** *maximum-computed-segment-lists* option under the **compute-profile** configuration statement. This configuration determines the maximum number of such segment lists computed for a given primary and secondary LSP.

- **Maximum segment list depth**

The maximum segment list depth computation parameter ensures that amongst the ECMP paths that satisfy all other constraints such as administrative group, only the paths that have segment lists less than or equal to the maximum segment list depth are used. When you configure this parameter as a constraint under the compute-profile, it overrides the **maximum-segment-list-depth** configuration under the **[edit protocols source-packet-routing]** hierarchy level, if present.

You can configure this attribute using the **maximum-segment-list-depth** *maximum-segment-list-depth* option under the **compute-profile** configuration statement.

- **Protected or unprotected adjacency SIDs**

You can configure protected or unprotected adjacency SID as a constraint under the **compute-profile** to avoid links with the specified SID type.

- **Metric type**

You can specify the type of metric on the link to be used for computation. By default, SR-TE LSPs use traffic-engineering metrics of the links for computation. The traffic-engineering metric for links is

advertised by traffic-engineering extensions of IGP protocols. However, you may also choose to use the IGP-metric for computation by using the metric-type configuration in the compute profile.

You can configure this attribute using the **metric-type (igp | te)** option under the **compute-profile** configuration statement.

Distributed CSPF Computation

The SRTE candidate paths are computed locally such that they satisfy the configured constraints. When label stack compression is disabled, the multi-path CSPF computation result is a set of adjacency SID stacks. When label stack compression is enabled, the result is a set of compressed label stacks (composed of adjacent SIDs and node SIDs).

When secondary paths are computed, the links, nodes and SRLGs taken by the primary paths are not avoided for computation. For more information on primary and secondary paths, see [“Configuring Primary and Secondary LSPs” on page 565](#).

For any LSPs with unsuccessful computation result, the computation is retried as traffic-engineering database (TED) changes.

Interaction Between Distributed CSPF Computation and SRTE Features

IN THIS SECTION

- [Weights Associated With Paths of an SRTE Policy | 619](#)
- [BFD Liveliness Detection | 619](#)
- [inherit-label-nexthops | 620](#)
- [Auto-Translate Feature | 620](#)

Weights Associated With Paths of an SRTE Policy

You can configure weights against computed and static SRTE paths, which contribute to the next hops of the route. However, a single path that has computation enabled can result in multiple segment lists. These computed segment lists are treated as ECMP amongst themselves. You can assign hierarchical ECMP weights to these segments, considering the weights assigned to each of the configured primaries.

BFD Liveliness Detection

You can configure BFD liveliness detection for the computed primary or secondary paths. Every computed primary or secondary path can result in multiple segment lists, as a result, the BFD parameters configured against the segment lists are applied to all the computed segment lists. If all the active primary paths go down, the pre-programmed secondary path (if provided) becomes active.

inherit-label-nexthops

You are not required to explicitly enable the **inherit-label-nexthops** configuration under the **[edit protocols source-packet-routing segment-list segment-list-name]** hierarchy for the computed primary or secondary paths, as it is a default behavior.

Auto-Translate Feature

You can configure the auto-translate feature on the segment lists, and the primary or secondary paths with the auto-translate feature reference these segment lists. On the other hand, the primary or secondary on which compute feature is enabled cannot reference any segment list. As a result, you cannot enable both the compute feature and the auto-translate feature for a given primary or secondary path. However, you could have an LSP configured with a primary path with compute type and another with auto-translate type.

Distributed CSPF Computation Sample Configurations

IN THIS SECTION

- [Example 1 | 620](#)
- [Example 2 | 621](#)
- [Example 3 | 621](#)

Example 1

In Example 1,

- The non-computed primary path references a configured segment-list. In this example, the configured segment list *static_sl1* is referenced, and it also serves as the name for this primary path.
- A computed primary should have a name configured, and this name should not reference any configured segment list. In this example, *compute_segment1* is not a configured segment list.
- The *compute_profile_red* compute-profile is applied to the primary path with the name *compute_segment1*.
- The *compute_profile_red* compute-profile includes a segment list of type **compute**, which is used to specify the explicit path constraint for the computation.

```
[edit protocols source-packet-routing]
segment-list static_sl1{
  hop1 label 80000
}
segment-list exp_path1 {
```



```
hop1 ip-address 10.1.1.1 loose
hop2 ip-address 2.2.2.2
compute
}
compute-profile compute_profile_red {
  include-any red
  segment-list exp_path1
  maximum-segment-list-depth 5
}
```

The weights for computed path next-hops and static next-hops are 2 and 3, respectively. Assuming the next-hops for computed paths are *comp_nh1*, *comp_nh2*, and *comp_nh3*, and the next-hop for static path is *static_nh*, the weights are applied as follows:

Next-Hop	Weight
comp_nh1	2
comp_nh2	2
comp_nh3	2
static_nh	9

Example 2

In Example 2, both the primary and secondary paths can be of compute type and can have their own compute-profiles.

```
[edit protocols source-packet-routing]
compute-profile compute_profile_green{
  include-any green
  maximum-segment-list-depth 5
}
compute-profile compute_profile_red{
  include-any red
  maximum-segment-list-depth 8
}
```

Example 3

In Example 3, when compute is mentioned under a primary or secondary path, it results in local computation of a path to the destination without any constraints or other parameters for the computation.


```
[edit protocols source-packet-routing]
source-routing-path srte_colored_policy1 {
  to 5.5.5.5
  color 5
  binding-sid 10001
  primary {
    compute_segment1 {
      compute
    }
  }
}
```

RELATED DOCUMENTATION

[compute-profile](#) | 2575

[compute-options](#) | 2574

Color-Based Mapping of VPN Services Overview

IN THIS SECTION

- [VPN Service Coloring](#) | 623
- [Specifying VPN Service Mapping Mode](#) | 626
- [Color-IP Protocol Next Hop Resolution](#) | 627
- [Fallback to IP Protocol Next Hop Resolution](#) | 628
- [Supported and Unsupported Features for Color-Based Mapping of VPN Services](#) | 628

You can specify color as a protocol next hop constraint (in addition to the IPv4 or IPv6 address) for resolving transport tunnels over static colored and BGP segment routing traffic-engineered (SRTE) LSPs. This is called the color-IP protocol next hop resolution, where you are required to configure a resolution-map and apply to the VPN services. With this feature, you can enable color-based traffic steering of Layer 2 and Layer 3 VPN services.

Junos OS supports colored SRTE LSPs associated with a single color. The color-based mapping of VPN services feature is supported on static colored LSPs and BGP SRTE LSPs.

VPN Service Coloring

IN THIS SECTION

- [Egress Color Assignment | 623](#)
- [Ingress Color Assignment | 625](#)

In general, a VPN service may be assigned a color on the egress router where the VPN NLRI is advertised, or on an ingress router where the VPN NLRI is received and processed.

You can assign a color to the VPN services at different levels:

- Per routing instance.
- Per BGP group.
- Per BGP neighbor.
- Per prefix.
- Set of prefixes.

Once you assign a color, the color is attached to a VPN service in the form of BGP color extended community.

You can assign multiple colors to a VPN service, referred to as multi-color VPN services. In such cases, the smallest color value is considered as the color of the VPN service, and all other colors are ignored.

Multiple colors are assigned by egress and/or ingress devices through multiple policies in the following order:

- BGP export policy on the egress device.
- BGP import policy on the ingress device.
- VRF import policy on the ingress device.

The two modes of VPN service coloring are:

Egress Color Assignment

In this mode, the egress device (that is, the advertiser of the VPN NLRI) is responsible for coloring the VPN service. To enable this mode, you can define a routing policy, and apply it in the VPN service's routing-instance **vrf-export**, group export, or group neighbor export at the **[edit protocols bgp]** hierarchy level. The VPN NLRI is advertised by BGP with the specified color extended community.

For example:


```
[edit routing-options]
community red-comm {
  members color:0:50;
}
```

```
[edit policy-options]
policy-statement pol-color {
  term t1 {
    from {
      [any match conditions];
    }
    then {
      community add red-comm;
      accept;
    }
  }
}
```

```
[edit routing-instances]
vpn-X {
  ...
  vrf-export pol-color ...;
}
```

Or

NOTE: When you apply the routing policy as an export policy of a BGP group or BGP neighbor, you must include the **vpn-apply-export** statement at the BGP, BGP group, or BGP neighbor level in order for the policy to take an effect on the VPN NLRI.

```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-A {
    export pol-color ...;
    vpn-apply-export;
  }
}
```


The routing policies are applied to Layer 3 VPN prefix NLRI, Layer 2 VPN NLRI, and EVPN NLRI. The color extended community is inherited by all the VPN routes, imported, and installed in the target VRFs on one or multiple ingress devices.

Ingress Color Assignment

In this mode, the ingress device (that is, the receiver of the VPN NLRI) is responsible for coloring the VPN service. To enable this mode, you can define a routing policy, and apply it to the VPN service's routing-instance **vrf-import**, group import, or group neighbor import at the **[edit protocols bgp]** hierarchy level. All the VPN routes matching the routing policy is attached with the specified color extended community.

For example:

```
[edit routing-options]
community red-comm {
  members color:0:50;
}
```

```
[edit policy-options]
policy-statement pol-color {
  term t1 {
    from {
      [any match conditions];
    }
    then {
      community add red-comm;
      accept;
    }
  }
}
```

```
[edit routing-instances]
vpn-Y {
  ...
  vrf-import pol-color ...;
}
```

Or

```
[edit protocols bgp]
group PEs {
  ...
}
```



```

neighbor PE-B {
    import pol-color ...;
}

```

Specifying VPN Service Mapping Mode

To specify flexible VPN service mapping modes, you must define a policy using the **resolution-map** statement, and refer the policy in a VPN service's routing-instance **vrf-import**, group import, or group neighbor import at the **[edit protocols bgp]** hierarchy level. All the VPN routes matching the routing policy are attached with the specified resolution-map.

For example:

```

[edit policy-options]
resolution-map map-A {
    <mode-1>;
    <mode-2>;
    ...
}
policy-statement pol-resolution {
    term t1 {
        from {
            [any match conditions];
        }
        then {
            resolution-map map-A;
            accept;
        }
    }
}

```

You can apply import policy to the VPN service's routing-instance.

```

[edit routing-instances]
vpn-Y {
    ...
    vrf-import pol-resolution ...;
}

```

You can also apply the import policy to a BGP group or BGP neighbor.


```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-B {
    import pol-resolution ...;
  }
}
```

NOTE: Each VPN service mapping mode should have a unique name defined in the resolution-map. Only a single entry of IP-color is supported in the resolution-map, where the VPN route(s) are resolved using a colored-IP protocol next hop in the form of **ip-address:color** over the inetcolor.0 and inet6color.0 routing tables.

Color-IP Protocol Next Hop Resolution

The protocol next hop resolution process is enhanced to support colored-IP protocol next hop resolution. For a colored VPN service, the protocol next hop resolution process takes a color and a resolution-map, builds a colored-IP protocol next hop in the form of *ip-address:color*, and resolves the protocol next hop in the inetcolor.0 and inet6color.0 routing tables.

You must configure a policy to support multipath resolution of colored Layer 2 VPN, Layer 3 VPN, or EVPN services over colored LSPs. The policy must then be applied with the relevant RIB table as the resolver import policy.

For example:

```
[edit policy-options]
policy-statement mpath {
  then multipath-resolve;
}
```

```
[edit routing-options]
resolution {
  rib bgp.l3vpn.0 {
    inetcolor-import mpath;
  }
}
resolution {
  rib bgp.l3vpn-inet6.0 {
```



```

        inet6color-import mpath;
    }
}
resolution {
    rib bgp.l2vpn.0 {
        inetcolor-import mpath;
    }
}
resolution {
    rib mpls.0 {
        inetcolor-import mpath;
    }
}
resolution {
    rib bgp.evpn.0 {
        inetcolor-import mpath;
    }
}

```

Fallback to IP Protocol Next Hop Resolution

If a colored VPN service does not have a resolution-map applied to it, the VPN service ignores its color and falls back to the IP protocol next hop resolution. Conversely, if a non-colored VPN service has a resolution-map applied to it, the resolution-map is ignored, and the VPN service uses the IP protocol next hop resolution.

The fallback is a simple process from colored SRTE LSPs to LDP LSPs, by using a RIB group for LDP to install routes in inet{6}color.0 routing tables. A longest prefix match for a colored-IP protocol next hop ensures that if a colored SRTE LSP route does not exist, an LDP route with a matching IP address should be returned.

Supported and Unsupported Features for Color-Based Mapping of VPN Services

The following features and functionality are supported with color-based mapping of VPN services:

- BGP Layer 3 VPN
- BGP Layer 2 VPN (Kompella Layer 2 VPN)
- BGP EVPN
- Resolution-map with a single IP-color option.
- Colored IPv4 and IPv6 protocol next hop resolution.

- Routing information base (also known as routing table) group based fallback to LDP LSP in inetcolor.0 or inet6color.0 routing tables.
- Colored SRTE LSP.
- Virtual platforms.
- 64-bit Junos OS.
- Logical systems.

The following features and functionality are not supported with color-based mapping of VPN services:

- Colored MPLS LSPs, such as RSVP, LDP, BGP-LU, static.
- Layer 2 circuit
- FEC-129 BGP auto-discovered and LDP-signaled Layer 2 VPN.
- VPLS
- MVPN
- IPv4 and IPv6 using resolution-map.
- BGP labeled unicast.

RELATED DOCUMENTATION

[Understanding Static Segment Routing LSP in MPLS Networks | 632](#)
[resolution-map | 2758](#)

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, liveness 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 83](#).

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 88](#).

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 | 630](#)
2. [Configuring the Provider and the Egress PE Switch | 631](#)

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 | 83](#)

[Configuring MPLS on Provider Switches | 88](#)

[Understanding MPLS Label Operations | 422](#)

Static Segment Routing Label Switched Path

IN THIS SECTION

- [Understanding Static Segment Routing LSP in MPLS Networks | 632](#)
- [Example: Configuring Static Segment Routing Label Switched Path | 657](#)

The segment routing architecture enables the ingress devices in a core network to steer traffic through explicit paths. You can configure these paths using segment lists to define the paths that the incoming traffic should take. The incoming traffic may be labeled or IP traffic, causing the forwarding operation at the ingress device to be either a label swap, or a destination-based lookup.

Understanding Static Segment Routing LSP in MPLS Networks

IN THIS SECTION

- [Introduction to Segment Routing LSPs | 633](#)
- [Benefits of using Segment Routing LSPs | 634](#)
- [Colored Static Segment Routing LSP | 634](#)
- [Non-Colored Static Segment Routing LSP | 635](#)
- [Static Segment Routing LSP Provisioning | 642](#)
- [Static Segment Routing LSP Limitations | 642](#)
- [Dynamic Creation of Segment Routing LSPs | 643](#)
- [Color-Based Mapping of VPN Services | 649](#)
- [Tunnel Templates for PCE-Initiated Segment Routing LSPs | 656](#)

Source packet routing or segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the actual path it should take.

Introduction to Segment Routing LSPs

Segment routing leverages the source routing paradigm. A device steers a packet through an ordered list of instructions, called segments. A segment can represent any instruction, topological or service-based. A segment can have a local semantic to a segment routing node or to a global node within a segment routing domain. Segment routing enforces a flow through any topological path and service chain while maintaining per-flow state only at the ingress device to the segment routing domain. Segment routing can be directly applied to the MPLS architecture with no change on the forwarding plane. A segment is encoded as an MPLS label. An ordered list of segments is encoded as a stack of labels. The segment to process is on the top of the stack. Upon completion of a segment, the related label is popped from the stack.

Segment routing LSPs can either be dynamic or static in nature.

Dynamic segment routing LSPs—When a segment routing LSP is created either by an external controller and downloaded to an ingress device through Path Computation Element Protocol (PCEP) extensions, or from a BGP segment routing policy through BGP segment routing extensions, the LSP is dynamically provisioned. The segment list of the dynamic segment routing LSP is contained in the PCEP Explicit Route Object (ERO), or the BGP segment routing policy of the LSP.

Static segment routing LSPs—When a segment routing LSP is created on the ingress device through local configuration, the LSP is statically provisioned.

A static segment routing LSP can further be classified as colored and non-colored LSPs based on the configuration of the **color** statement at the **[edit protocols source-packet-routing source-routing-path *lsp-name*]** hierarchy level.

For example:

```
[edit protocols]
  source-packet-routing {
    source-routing-path lsp_name {
      to destination_address;
      color color_value;
      binding-sid binding-label;
      primary segment_list_1_name weight weight;
      ...
      primary segment_list_n_name weight weight;
      secondary segment_list_n_name;
      sr-preference sr_preference_value;
    }
  }
```

Here, each primary and secondary statement refers to a segment list.

```
[edit protocols]
  source-packet-routing {
    segment-list segment_list_name {
      hop_1_name label sid_label;
      ...
      hop_n_name label sid_label;
    }
  }
```

Benefits of using Segment Routing LSPs

- Static segment routing does not rely on per LSP forwarding state on transit routers. Hence, removing the need of provisioning and maintaining per LSP forwarding state in the core.
- Provide higher scalability to MPLS networks.

Colored Static Segment Routing LSP

IN THIS SECTION

- [Understanding Colored Static Segment Routing LSPs | 635](#)
- [Segment List of Colored Segment Routing LSPs | 635](#)

A static segment routing LSP configured with the **color** statement is called a colored LSP.

Understanding Colored Static Segment Routing LSPs

Similar to a BGP segment routing policy, the ingress route of the colored LSP is installed in the **inetcolor.0** or **inet6color.0** routing tables, with **destination-ip-address, color** as key for mapping IP traffic.

A static colored segment routing LSP may have a binding SID, for which a route is installed in the **mpls.0** routing table. This binding SID label is used to map labeled traffic to the segment routing LSP. The gateways of the route are derived from the segment list configurations under the primary and secondary paths.

Segment List of Colored Segment Routing LSPs

The colored static segment routing LSPs already provide support for first hop label mode of resolving an LSP. However, first hop IP mode is not supported for colored segment routing LSPs. Starting in Junos OS Release 19.1R1, a commit check feature is introduced to ensure that all the segment lists contributing to the colored routes have the minimum label present for all hops. If this requirement is not met, the commit is blocked.

Non-Colored Static Segment Routing LSP

IN THIS SECTION

- [Understanding Non-Colored Segment Routing LSPs | 635](#)
- [Segment List of Non-Colored Segment Routing LSPs | 636](#)

A static segment routing LSP that is configured without the **color** statement is a non-colored LSP. Similar to PCEP segment routing tunnels, the ingress route is installed in the **inet.3** or **inet6.3** routing tables.

Junos OS supports non-colored static segment routing LSPs on ingress routers. You can provision non-colored static segment routing LSP by configuring one source routed path and one or more segment lists. These segment lists can be used by multiple non-colored segment routing LSPs.

Understanding Non-Colored Segment Routing LSPs

The non-colored segment routing LSP has a unique name and a destination IP address. An ingress route to the destination is installed in the **inet.3** routing table with a default preference of 8 and a metric of 1. This route allows non-colored services to be mapped to the segment routing LSP pertaining to the destination. In case the non-colored segment routing LSP does not require an ingress route then the ingress route can be disabled. A non-colored segment routing LSP uses binding SID label to achieve segment routing LSP stitching. This label that can be used to model the segment routing LSP as a segment that may be further used to construct other segment routing LSPs in a hierarchical manner. The transit of the binding SID label, by default, has a preference of 8 and a metric of 1.

Starting in Junos OS Release 18.2R1, statically configured non-colored segment routing LSPs on the ingress device are reported to the Path Computation Element (PCE) through a Path Computation Element Protocol

(PCEP) session. These non-colored segment routing LSPs may have binding service identifier (SID) labels associated with them. With this feature, the PCE can use this binding SID label in the label stack to provision PCE-initiated segment routing LSP paths.

A non-colored segment routing LSP can have a maximum of 8 primary paths. If there are multiple operational primary paths then the packet forwarding engine (PFE) distributes traffic over the paths based on the load balancing factors like the weight configured on the path. This is equal cost multi path (ECMP) if none of the paths have a weight configured on them or weighted ECMP if at least one of the paths has a non-zero weight configured on the paths. In both the cases, when one or some of the paths fail, the PFE rebalances the traffic over the remaining paths that automatically leads to achieving path protection. A non-colored segment routing LSP can have a secondary path for dedicated path protection. Upon failure of a primary path, the PFE rebalances the traffic to the remaining functional primary paths. Otherwise, the PFE switches the traffic to the backup path, hence achieving path protection. A non-colored segment routing LSP may specify a metric at `[edit protocols source-packet-routing source-routing-path lsp-name]` for its ingress and binding-SID routes. Multiple non-colored segment routing LSPs have the same destination address that contribute to the next hop of the ingress route.

Multiple non-colored segment routing LSPs have the same destination address that contribute to the next hop of the ingress route. Each path, either primary or secondary, of each segment routing LSP is considered as a gateway candidate, if the path is functional and the segment routing LSP has the best preference of all these segment routing LSPs. However, the maximum number of gateways that the next-hop can hold cannot exceed the RPD multi-path limit, which is 128 by default. Extra paths are pruned, firstly secondary paths and then primary paths. A given segment list may be referred multiple times as primary or secondary paths by these segment routing LSPs. In this case, there are multiple gateways, each having a unique segment routing LSP tunnel ID. These gateways are distinct, although they have identical outgoing label stack and interface. A non-colored segment routing LSP and a colored segment routing LSP may also have the same destination address. However, they correspond to different destination addresses for ingress routes, as the colored segment routing LSP's destination address is constructed with both its destination address and color.

NOTE: In the case where a static non-colored segment routing LSP and a PCEP-created segment routing LSP co-exist and have the same address that contributes to the same ingress route, if they also have the same preference. Otherwise, the segment routing LSP with the best preference is installed for the route.

Segment List of Non-Colored Segment Routing LSPs

A segment list consists of a list of hops. These hops are based on the SID label or an IP address. The number of SID labels in the segment list should not exceed the maximum segment list limit. You can configure the maximum segment list limit at the `[edit protocols source-packet-routing]` hierarchy level.

Prior to Junos OS Release 19.1R1, for a non-colored static segment routing LSP to be usable, the first hop of the segment list had to be an IP address of an outgoing interface and the second to *n*th hops could be

SID labels. Starting in Junos OS Release 19.1R1, this requirement does not apply, as the first hop of the non-colored static LSPs now provides support for SID labels, in addition to IP addresses. With the first hop label support, MPLS fast reroute (FRR) and weighted equal-cost multipath is enabled for resolving the static non-colored segment routing LSPs, similar to colored static LSPs.

For the first-hop label mode to take effect, you must include the **inherit-label-nexthops** statement globally or individually for a segment list, and the first hop of the segment list must include both IP address and label. If the first hop includes only IP address, the **inherit-label-nexthops** statement does not have any effect.

You can configure **inherit-label-nexthops** at any one of the following hierarchies. The **inherit-label-nexthops** statement takes effect only if the segment list first hop includes both IP address and label.

- **Segment list level**—At the `[edit protocols source-packet-routing segment-list segment-list-name]` hierarchy level.
- **Globally**—At the `[edit protocols source-packet-routing]` hierarchy level.

When the **inherit-label-nexthops** statement is configured globally, it takes precedence over the segment-list level configuration, and the **inherit-label-nexthops** configuration is applied to all the segment lists. When the **inherit-label-nexthops** statement is not configured globally, only segment lists with both labels and IP address present in the first hop, and configured with **inherit-label-nexthops** statement are resolved using SID labels.

For dynamic non-colored static LSPs, that is the PCEP-driven segment routing LSPs, the **inherit-label-nexthops** statement must be enabled globally, as the segment-level configuration is not applied.

[Table 16 on page 637](#) describes the mode of segment routing LSP resolution based on the first hop specification.

Table 16: Non-Colored Static LSP Resolution Based on First Hop Specification

First Hop Specification	Mode of LSP Resolution
IP address only For example: <pre>segment-list path-1 { hop-1 ip-address 172.0.12.2; hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	The segment list is resolved using the IP address.

Table 16: Non-Colored Static LSP Resolution Based on First Hop Specification (*continued*)

First Hop Specification	Mode of LSP Resolution
<p>SID only</p> <p>For example:</p> <pre>segment-list path-2 { hop-1 label 1000011; hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	<p>The segment list is resolved using SID labels.</p>
<p>IP address and SID (without the inherit-label-nexthops configuration)</p> <p>For example:</p> <pre>segment-list path-3 { hop1 { label 801006; ip-address 172.24.1.2; } hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	<p>By default, the segment list is resolved using IP address.</p>
<p>IP address and SID (with the inherit-label-nexthops configuration)</p> <p>For example:</p> <pre>segment-list path-3 { inherit-label-nexthops; hop1 { label 801006; ip-address 172.24.1.2; } hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	<p>The segment list is resolved using SID labels.</p>

You can use the **show route *ip-address* protocol spring-te active-path table inet.3** command to view the non-colored segment routing traffic-engineered LSPs having multiple segment lists installed in the inet.3 routing table.

For example:

```
user@host> show route 7.7.7.7 protocol spring-te active-path table inet.3
```

```
inet.3: 42 destinations, 59 routes (41 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

7.7.7.7/32      *[SPRING-TE/8] 00:01:25, metric 1, metric2 0
                > to 11.1.1.2 via et-0/0/0.1, Push 801007
                  to 21.1.1.2 via et-0/0/2.1, Push 801007
                  to 11.102.1.2 via et-0/0/0.2, Push 801007, Push 801002(top)

                  to 21.202.1.2 via et-0/0/2.2, Push 801007, Push 801005(top)

                  to 11.103.1.2 via et-0/0/0.3, Push 801007, Push 801003(top)

                  to 21.203.1.2 via et-0/0/2.3, Push 801007, Push 801006(top)

                  to 11.104.1.2 via et-0/0/0.4, Push 801007, Push 801003, Push
801002(top)
                  to 21.204.1.2 via et-0/0/2.4, Push 801007, Push 801006, Push
801005(top)
```


NOTE:

The first hop type of segment lists of a static segment routing LSP can cause a commit to fail, if:

- Different segment lists of a tunnel have different first hop resolution types. This is applicable to both colored and non-colored static segment routing LSPs. However, this does not apply for PCEP-driven LSPs; a system log message is generated for the mismatch in the first hop resolution type at the time of computing the path.

For example:

```
segment-list path-1 {
  hop-1 ip-address 172.0.12.2;
  hop-2 label 1000012;
  hop-3 label 1000013;
  hop-4 label 1000014;
}
segment-list path-2 {
  hop-1 label 1000011;
  hop-2 label 1000012;
  hop-3 label 1000013;
  hop-4 label 1000014;
}
source-routing-path lsp1 {
  to 172.10.10.1;
  primary {
    path-1;
    path-2;
  }
}
```

The commit of tunnel *lsp1* fails, as path-1 is of IP address mode and path-2 is of label mode.

- The binding SID is enabled for the static non-colored LSP whose segment list type is SID label.

For example:

```
segment-list path-3 {
  hop-1 label 1000011;
  hop-2 label 1000012;
  hop-3 label 1000013;
  hop-4 label 1000014;
}
source-routing-path lsp1 {
```



```

to 172.10.10.1;
binding-sid 333;
primary {
    path-3;
}
}

```

Configuring binding SID over label segment list is supported only for colored static LSPs and not for no-colored static LSPs.

Static Segment Routing LSP Provisioning

Segment provisioning is performed on per-router basis. For a given segment on a router, a unique service identifier (SID) label is allocated from a desired label pool which may be from the dynamic label pool for an adjacency SID label or from the segment routing global block (SRGB) for a prefix SID or node SID. The adjacency SID label can be dynamically allocated, which is the default behavior, or be allocated from a local static label pool (SRLB). A route for the SID label is then installed in the `mpls.0` table.

Junos OS allows static segment routing LSPs by configuring the **segment** statement at the `[edit protocols mpls static-label-switched-path static-label-switched-path]` hierarchy level. A static segment LSP is identified by a unique SID label that falls under Junos OS static label pool. You can configure the Junos OS static label pool by configuring the **static-label-range static-label-range** statement at the `[edit protocols mpls label-range]` hierarchy level.

Static Segment Routing LSP Limitations

- Junos OS currently has a limitation that the next hop cannot be built to push more than the maximum segment list depth labels. So, a segment list with more than the maximum SID labels (excluding the SID label of the first hop which is used to resolve forwarding next-hop) is not usable for colored or non-colored segment routing LSPs. Also, the actual number allowed for a given segment routing LSP may be even lower than the maximum limit, if an MPLS service is on the segment routing LSP or the segment routing LSP is on a link or a node protection path. In all cases, the total number of service labels, SID labels, and link or node protection labels must not exceed the maximum segment list depth. You can configure the maximum segment list limit at `[edit protocols source-packet-routing]` hierarchy level. Multiple non-colored segment routing LSPs with less than or equal to the maximum SID labels can be stitched together to construct a longer segment routing LSP. This is called segment routing LSP stitching. It can be achieved using binding-SID label.
- The segment routing LSP stitching is actually performed at path level. If a non-colored segment routing LSP has multiple paths that is multiple segment lists, each path can be independently stitched to another non-colored segment routing LSP at a stitching point. A non-colored segment routing LSP which is dedicated to stitching may disable ingress route installation by configuring **no-ingress** statement at `[edit protocols source-packet-routing source-routing-path lsp-name]` hierarchy level.

- A maximum of 8 primary paths and 1 secondary path are supported per non-colored static segment routing LSP. If there is a violation in configuration, commit check fails with an error.
- If any segment-list is configured with more labels than the maximum segment list depth, the configuration commit check fails with an error.

Dynamic Creation of Segment Routing LSPs

IN THIS SECTION

- [Configuring Dynamic Segment Routing LSP Template | 643](#)
- [Resolving Dynamic Segment Routing LSPs | 644](#)
- [Considerations for Configuring Dynamic Creation of Segment Routing LSPs | 648](#)
- [Services Supported over Dynamic Segment Routing LSPs | 648](#)
- [Behavior With Multiple Tunnel Sources in Segment Routing | 649](#)
- [Dynamic Segment Routing LSPs Limitations | 649](#)

In segment routing networks that have each provider edge (PE) device connected to every other PE device, a large amount of configuration is required for setting up the segment routing label-switched paths (LSPs), although only a few segment routing traffic-engineered (SR-TE) paths may be used. You can enable BGP-triggered dynamic creation of these LSPs to reduce the amount of configuration in such deployments.

Configuring Dynamic Segment Routing LSP Template

To configure the template for enabling dynamic creation of segment routing LSPs, you must include the [spring-te](#) statement at the `[edit routing-options dynamic-tunnels]` hierarchy.

- The following is a sample configuration for color dynamic segment routing LSP template:

```
[edit routing-options]
dynamic-tunnels {
  <dynamic-tunnel-name> {
    spring-te {
      source-routing-path-template {
        <template-name1> color [c1 c2];
        <template-name2> color [c3];
        <template-name3> color-any;
      }
      destination-networks {
        <dest1>;
        <dest2>;
      }
    }
  }
}
```



```

    }
  }
}

```

- The following is a sample configuration for non-color dynamic segment routing LSP template:

```

dynamic-tunnels {
  <dynamic-tunnel-name> {
    spring-te {
      source-routing-path-template {
        <template-name1>;
      }
      destination-networks {
        <dest1>;
        <dest2>;
      }
    }
  }
}

```

Resolving Dynamic Segment Routing LSPs

IN THIS SECTION

- [Resolving Colored Dynamic Segment Routing LSP | 644](#)
- [Resolving Uncolored Dynamic Segment Routing LSPs | 645](#)
- [Dynamic Segment Routing LSP Sample Configuration | 645](#)

Resolving Colored Dynamic Segment Routing LSP

When the BGP prefixes are assigned with color community, they initially get resolved over the catch-all-route-for-that-particular-color policy, and in turn, the SR-TE template on which the BGP prefix should be resolved onto is identified. The destinations SID is then derived from the BGP payload prefix next-hop attribute. For example, if the next hop of the BGP payload prefix is an IP address that belongs to Device A, then the node-SID of Device A is taken and a corresponding label is prepared and pushed to the bottom of the stack. The BGP payload prefix is resolved in a color-only mode, where the node-SID of Device A is at the bottom of the final label stack, and the SR-TE path labels are on top.

The final LSP template name is a combination of prefix, color, and tunnel name; for example, **<prefix>:<color>:dt-srte-<tunnel-name>**. The color in the LSP name is displayed in hexadecimal format, and the format of the tunnel name is similar to that of RSVP-triggered tunnel LSP names.

To successfully resolve a colored destination network, the color should have a valid template mapping, either to a specific color, or through the **color-any** template. Without a valid mapping, the tunnel is not created and the BGP route requesting for resolution remains unresolved.

Resolving Uncolored Dynamic Segment Routing LSPs

The catch-all routes for non-colored LSPs are added to the inet.3 routing table. The non-colored tunnel destination must be configured in a different **spring-te** configuration with only one template name in the mapping list. This template name is used for all the tunnel routes matching any of the destination networks configured under the same **spring-te** configuration. These tunnels are similar to RSVP tunnels in functionality.

The final LSP template name is a combination of prefix and tunnel name; for example, **<prefix>:dt-srte-<tunnel-name>**.

Dynamic Segment Routing LSP Sample Configuration

IN THIS SECTION

- [Colored Dynamic Segment Routing LSPs | 645](#)
- [Non-Colored Dynamic Segment Routing LSPs | 646](#)
- [Unresolved Dynamic Segment Routing LSP | 647](#)

The dynamic segment routing LSP template always carries a partial path. The last hop node SID is derived automatically at the tunnel creation time depending on the protocol next-hop address (PNH) node SID. The same template can be used by multiple tunnels to different destinations. In such cases, the partial path remains the same, and only the last hop changes depending on the PNH. Dynamic segment routing LSP templates are not common to a single tunnel, as a result a full path cannot be carried on it. You can use a segment list if a full path is to be used.

Colored Dynamic Segment Routing LSPs

Sample configuration for colored dynamic segment routing LSPs:

```
protocols source-packet-routing {
  source-routing-path-template sr_lsp1 {
    primary sr_sl1
    primary sr_sl2 weight 2
    sr-preference 180;
  }
}
dynamic-tunnels tunnel1 {
  spring-te {
    source-routing-path-template {
```



```

        sr_lsp1 color [ 123 124 125 ];
        sr_lsp2 color-any
    }
    destination-networks {
        22.33.44.0/24;
    }
}
}
}

```

For the above-mentioned sample configuration, the route entries are as follows:

inetcolor.0 tunnel route: 22.33.44.0-0/24 --> RT_NH_TUNNEL

inetcolor6.0 tunnel route: ffff::22.33.44.0-0/120 --> RT_NH_TUNNEL

BGP prefix to tunnel mapping:

R1(prefix) -> 22.33.44.55-101(PNH) LSP tunnel name = 22.33.44.55:65:dt-srte-tunnel1

R1(prefix) -> ffff::22.33.44.55-101(PNH) LSP tunnel name = 22.33.44.55:65:dt-srte-tunnel1

R1(prefix) -> ffff::22.33.44.55-124(PNH) LSP tunnel name = 22.33.44.55:7c:dt-srte-tunnel1

inetcolor.0 tunnel route:

22.33.44.55-101/64 --> <next-hop>

22.33.44.55-124/64 --> <next-hop>

inetcolor6.0 tunnel route:

ffff::22.33.44.55-101/160 --> <next-hop>

ffff::22.33.44.55-124/160 --> <next-hop>

Non-Colored Dynamic Segment Routing LSPs

Sample configuration for non-colored dynamic segment routing LSPs:

```

protocols source-packet-routing {
    source-routing-path-template sr_lsp1 {
        primary sr_sl1
        primary sr_sl2 weight 2
        sr-preference 180;
    }
}
dynamic-tunnels {
    tunnel1 {
        spring-te {
            source-routing-path-template {

```



```

source-routing-path-template {
    sr_lsp1 color [120 121 122 123];
}
destination-networks {
    22.33.44.0/24;
    1.1.1.0/24;
}
}
}

```

For the above-mentioned sample configuration, the route entries are as follows:

inetcolor.0 tunnel route: 22.33.44.0 - 0/24 --> RT_NH_TUNNEL 1.1.1.0 - 0/24 --> RT_NH_TUNNEL

inetcolor6.0 tunnel route: ffff::22.33.44.0 - 0/120 --> RT_NH_TUNNEL ffff::1.1.1.0 - 0/24 --> RT_NH_TUNNEL

BGP prefix to tunnel mapping: R1(prefix) -> 22.33.44.55-124(PNH) Tunnel will not be created. (Template not found for the color).

Considerations for Configuring Dynamic Creation of Segment Routing LSPs

When configuring the dynamic creation of segment routing LSPs, take the following into consideration:

- A template can be assigned with a color object. When the dynamic tunnel **spring-te** configuration includes a template with a color object, you must configure all other templates with color objects as well. All destinations are assumed to be colored within that configuration.
- A template can have a list of colors defined on it, or can be configured with the **color-any** option. Both these options can coexist in the same **spring-te** configuration. In such cases, templates assigned with specific colors have a higher preference.
- In a **spring-te** configuration, only one template can be defined with the **color-any** option.
- The color-to-template mapping is done on a one-to-one basis. One color cannot map to multiple templates.
- The template name should be configured in the **spring-te** statement under the **[edit protocols]** hierarchy, and should have the **primary** option enabled.
- Colored and non-colored destinations cannot co-exist in the same **spring-te** configuration.
- You cannot configure same destination networks, with or without color, under different **spring-te** configuration statements.
- In non-colored **spring-te** configuration, only one template can be configured without color object.

Services Supported over Dynamic Segment Routing LSPs

The following services are supported over colored dynamic segment routing LSPs:

- Layer 3 VPN

- BGP EVPN
- Export policy services

The following services are supported over non-colored dynamic segment routing LSPs:

- Layer 3 VPN
- Layer 2 VPN
- Multipath configurations

Behavior With Multiple Tunnel Sources in Segment Routing

When two sources download routes to the same destination from segment routing (for example static and dynamic sourced tunnels), then the segment routing preference is used for choosing the active route entry. A higher value has greater preference. In case the preference remains the same, then the tunnel source is used to determine the route entry.

Dynamic Segment Routing LSPs Limitations

The dynamic SR-TE LSPs do not support the following features and functionalities:

- IPv6 segment routing tunnels.
- Static tunnels.
- 6PE is not supported.
- Distributed CSPF.
- sBFD and LDP tunnelling is not supported for dynamic SR-TE LSPs and in a template.
- Install and B-SID routes in a template.

Color-Based Mapping of VPN Services

IN THIS SECTION

- [VPN Service Coloring | 650](#)
- [Specifying VPN Service Mapping Mode | 653](#)
- [Color-IP Protocol Next Hop Resolution | 654](#)
- [Fallback to IP Protocol Next Hop Resolution | 655](#)
- [Supported and Unsupported Features for Color-Based Mapping of VPN Services | 655](#)

You can specify color as a protocol next hop constraint (in addition to the IPv4 or IPv6 address) for resolving transport tunnels over static colored and BGP segment routing traffic-engineered (SRTE) LSPs. This is called the color-IP protocol next hop resolution, where you are required to configure a resolution-map

and apply to the VPN services. With this feature, you can enable color-based traffic steering of Layer 2 and Layer 3 VPN services.

Junos OS supports colored SRTE LSPs associated with a single color. The color-based mapping of VPN services feature is supported on static colored LSPs and BGP SRTE LSPs.

VPN Service Coloring

IN THIS SECTION

- [Egress Color Assignment | 650](#)
- [Ingress Color Assignment | 652](#)

In general, a VPN service may be assigned a color on the egress router where the VPN NLRI is advertised, or on an ingress router where the VPN NLRI is received and processed.

You can assign a color to the VPN services at different levels:

- Per routing instance.
- Per BGP group.
- Per BGP neighbor.
- Per prefix.

Once you assign a color, the color is attached to a VPN service in the form of BGP color extended community.

You can assign multiple colors to a VPN service, referred to as multi-color VPN services. In such cases, the last color attached is considered as the color of the VPN service, and all other colors are ignored.

Multiple colors are assigned by egress and/or ingress devices through multiple policies in the following order:

- BGP export policy on the egress device.
- BGP import policy on the ingress device.
- VRF import policy on the ingress device.

The two modes of VPN service coloring are:

Egress Color Assignment

In this mode, the egress device (that is, the advertiser of the VPN NLRI) is responsible for coloring the VPN service. To enable this mode, you can define a routing policy, and apply it in the VPN service's

routing-instance **vrf-export**, group export, or group neighbor export at the **[edit protocols bgp]** hierarchy level. The VPN NLRI is advertised by BGP with the specified color extended community.

For example:

```
[edit routing-options]
community red-comm {
  members color:0:50;
}
```

```
[edit policy-options]
policy-statement pol-color {
  term t1 {
    from {
      [any match conditions];
    }
    then {
      community add red-comm;
      accept;
    }
  }
}
```

```
[edit routing-instances]
vpn-X {
  ...
  vrf-export pol-color ...;
}
```

Or

NOTE: When you apply the routing policy as an export policy of a BGP group or BGP neighbor, you must include the **vpn-apply-export** statement at the BGP, BGP group, or BGP neighbor level in order for the policy to take an effect on the VPN NLRI.

```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-A {
    export pol-color ...;
```



```

        vpn-apply-export;
    }
}

```

The routing policies are applied to Layer 3 VPN prefix NLRI, Layer 2 VPN NRLIs, and EVPN NLRI. The color extended community is inherited by all the VPN routes, imported, and installed in the target VRFs on one or multiple ingress devices.

Ingress Color Assignment

In this mode, the ingress device (that is, the receiver of the VPN NLRI) is responsible for coloring the VPN service. To enable this mode, you can define a routing policy, and apply it to the VPN service's routing-instance **vrf-import**, group import, or group neighbor import at the **[edit protocols bgp]** hierarchy level. All the VPN routes matching the routing policy is attached with the specified color extended community.

For example:

```

[edit routing-options]
community red-comm {
    members color:0:50;
}

```

```

[edit policy-options]
policy-statement pol-color {
    term t1 {
        from {
            [any match conditions];
        }
        then {
            community add red-comm;
            accept;
        }
    }
}

```

```

[edit routing-instances]
vpn-Y {
    ...
    vrf-import pol-color ...;
}

```

Or


```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-B {
    import pol-color ...;
  }
}
```

Specifying VPN Service Mapping Mode

To specify flexible VPN service mapping modes, you must define a policy using the **resolution-map** statement, and refer the policy in a VPN service's routing-instance **vrf-import**, group import, or group neighbor import at the **[edit protocols bgp]** hierarchy level. All the VPN routes matching the routing policy are attached with the specified resolution-map.

For example:

```
[edit policy-options]
resolution-map map-A {
  <mode-1>;
  <mode-2>;
  ...
}
policy-statement pol-resolution {
  term t1 {
    from {
      [any match conditions];
    }
    then {
      resolution-map map-A;
      accept;
    }
  }
}
```

You can apply import policy to the VPN service's routing-instance.

```
[edit routing-instances]
vpn-Y {
  ...
  vrf-import pol-resolution ...;
}
```

You can also apply the import policy to a BGP group or BGP neighbor.


```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-B {
    import pol-resolution ...;
  }
}
```

NOTE: Each VPN service mapping mode should have a unique name defined in the resolution-map. Only a single entry of IP-color is supported in the resolution-map, where the VPN route(s) are resolved using a colored-IP protocol next hop in the form of **ip-address:color**.

Color-IP Protocol Next Hop Resolution

The protocol next hop resolution process is enhanced to support colored-IP protocol next hop resolution. For a colored VPN service, the protocol next hop resolution process takes a color and a resolution-map, builds a colored-IP protocol next hop in the form of *IP-address:color*, and resolves the protocol next hop in the inet6color.0 routing table.

You must configure a policy to support multipath resolution of colored Layer 2 VPN, Layer 3 VPN, or EVPN services over colored LSPs. The policy must then be applied with the relevant RIB table as the resolver import policy.

For example:

```
[edit policy-options]
policy-statement mpath {
  then multipath-resolve;
}
```

```
[edit routing-options]
resolution {
  rib bgp.l3vpn.0 {
    inetcolor-import mpath;
  }
}
resolution {
  rib bgp.l3vpn-inet6.0 {
    inet6color-import mpath;
  }
}
```



```

resolution {
    rib bgp.l2vpn.0 {
        inetcolor-import mpath;
    }
}
resolution {
    rib mpls.0 {
        inetcolor-import mpath;
    }
}
resolution {
    rib bgp.evpn.0 {
        inetcolor-import mpath;
    }
}

```

Fallback to IP Protocol Next Hop Resolution

If a colored VPN service does not have a resolution-map applied to it, the VPN service ignores its color and falls back to the IP protocol next hop resolution. Conversely, if a non-colored VPN service has a resolution-map applied to it, the resolution-map is ignored, and the VPN service uses the IP protocol next hop resolution.

The fallback is a simple process from colored SRTE LSPs to LDP LSPs, by using a RIB group for LDP to install routes in inet{6}color.0 routing tables. A longest prefix match for a colored-IP protocol next hop ensures that if a colored SRTE LSP route does not exist, an LDP route with a matching IP address should be returned.

Supported and Unsupported Features for Color-Based Mapping of VPN Services

The following features and functionality are supported with color-based mapping of VPN services:

- BGP Layer 3 VPN
- BGP Layer 2 VPN (Kompella Layer 2 VPN)
- BGP EVPN
- Resolution-map with a single IP-color option.
- Colored IPv4 and IPv6 protocol next hop resolution.
- Routing information base (also known as routing table) group based fallback to LDP LSP in inetcolor.0 routing table.
- Colored SRTE LSP.
- Virtual platforms.

- 64-bit Junos OS.
- Logical systems.

The following features and functionality are not supported with color-based mapping of VPN services:

- Colored MPLS LSPs, such as RSVP, LDP, BGP-LU, static.
- Layer 2 circuit
- FEC-129 BGP auto-discovered and LDP-signaled Layer 2 VPN.
- VPLS
- MVPN
- IPv4 and IPv6 using resolution-map.
- BGP labeled unicast.

Tunnel Templates for PCE-Initiated Segment Routing LSPs

You can configure a tunnel template for PCE-initiated segment routing LSPs to pass down two additional parameters for these LSPs - Bidirectional forwarding detection (BFD) and LDP tunneling.

When a PCE-Initiated segment routing LSP is being created, the LSP is checked against policy statements (if any) and if there is a match, the policy applies the configured template for that LSP. The template configuration is inherited only if it is not provided by the LSP source (PCEP); for example, metric.

To configure a template:

1. Include the [source-routing-path-template](#) statement at the **[edit protocols source-packet-routing]** hierarchy level. You can configure the additional BFD and LDP tunneling parameters here.
2. Include the [source-routing-path-template-map](#) statement at the **[edit protocols source-packet-routing]** hierarchy level to list the policy statements against which the PCE-initiated LSP should be checked.
3. Define a policy to list the LSPs on which the template should be applied.

The **from** statement can include either the LSP name or LSP regular expression using the **lsp** and **lsp-regex** match conditions. These options are mutually exclusive, so you can specify only one option at a given point in time.

The **then** statement must include the **sr-te-template** option with an accept action. This applies the template to the PCE-initiated LSP.

Take the following into consideration when configuring a template for PCE-initiated LSPs:

- Template configuration is not applicable to statically configured segment routing LSPs, or any other client's segment routing LSP.
- PCEP-provided configuration has precedence over template configuration.

- PCEP LSP does not inherit template segment-list configuration.

Example: Configuring Static Segment Routing Label Switched Path

IN THIS SECTION

- [Requirements | 657](#)
- [Overview | 657](#)
- [Configuration | 658](#)
- [Verification | 671](#)

This example shows how to configure static segment routing label switched paths (LSPs) in MPLS networks. This configuration helps to bring higher scalability to MPLS networks.

Requirements

This example uses the following hardware and software components:

- Seven MX Series 5G Universal Routing Platforms
- Junos OS Release 18.1 or later running on all the routers

Before you begin, be sure you configure the device interfaces.


Overview

Junos OS a set of explicit segment routing paths are configured on the ingress router of a non-colored static segment routing tunnel by configuring the **segment-list** statement at the **[edit protocols source-packet-routing]** hierarchy level. You can configure segment routing tunnel by configuring the **source-routing-path** statement at **[edit protocols source-packet-routing]** hierarchy level. The segment routing tunnel has a destination address and one or more primary paths and optionally secondary paths that refer to the segment list. Each segment list consists of a sequence of hops. For non-colored static segment routing tunnel, the first hop of the segment list specifies an immediate next hop IP address and the second to Nth hop specifies the segment identifies (SID) labels corresponding to the link or node which the path traverses. The route to the destination of the segment routing tunnel is installed in inet.3 table.

Topology

In this example, configure layer 3 VPN on the provider edge routers PE1 and PE5. Configure the MPLS protocol on all the routers. The segment routing tunnel is configured from router PE1 to router PE5 with a primary path configured on router PE1 and router PE5 . Router PE1 is also configured with secondary path for path protection. The transit routers PE2 to PE4 are configured with adjacency SID labels with label pop and an outgoing interface.

Figure 50: Static Segment Routing Label Switched Path



Configuration

IN THIS SECTION

- [Configuring Device PE1 | 662](#)
- [Configuring Device PE2 | 668](#)
- [Results | 669](#)

CLI Quick Configuration

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

PE1

```
set interfaces ge-0/0/0 unit 0 family inet address 10.10.12.1/24
set interfaces ge-0/0/0 unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/1 unit 0 family inet address 10.10.13.1/24
set interfaces ge-0/0/1 unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/5 unit 0 family inet address 10.10.17.1/24
set routing-options autonomous-system 65000
set routing-options forwarding-table export load-balance-policy
set routing-options forwarding-table chained-composite-next-hop ingress l3vpn
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols bgp group pe type internal
set protocols bgp group pe local-address 192.168.147.211
set protocols bgp group pe family inet-vpn unicast
set protocols bgp group pe neighbor 192.168.146.181
```



```

set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols source-packet-routing segment-list sl-15-primary hop-1 ip-address 10.10.13.3
set protocols source-packet-routing segment-list sl-15-primary hop-2 label 1000134
set protocols source-packet-routing segment-list sl-15-primary hop-3 label 1000145
set protocols source-packet-routing segment-list sl-15-backup hop-1 ip-address 10.10.12.2
set protocols source-packet-routing segment-list sl-15-backup hop-2 label 1000123
set protocols source-packet-routing segment-list sl-15-backup hop-3 label 1000134
set protocols source-packet-routing segment-list sl-15-backup hop-4 label 1000145
set protocols source-packet-routing source-routing-path lsp-15 to 192.168.146.181
set protocols source-packet-routing source-routing-path lsp-15 binding-sid 1000999
set protocols source-packet-routing source-routing-path lsp-15 primary sl-15-primary
set protocols source-packet-routing source-routing-path lsp-15 secondary sl-15-backup
set policy-options policy-statement VPN-A-export term a from protocol ospf
set policy-options policy-statement VPN-A-export term a from protocol direct
set policy-options policy-statement VPN-A-export term a then community add VPN-A
set policy-options policy-statement VPN-A-export term a then accept
set policy-options policy-statement VPN-A-export term b then reject
set policy-options policy-statement VPN-A-import term a from protocol bgp
set policy-options policy-statement VPN-A-import term a from community VPN-A
set policy-options policy-statement VPN-A-import term a then accept
set policy-options policy-statement VPN-A-import term b then reject
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf from route-filter 10.10.0.0/16 orlonger
set policy-options policy-statement bgp-to-ospf then accept
set policy-options policy-statement load-balance-policy then load-balance per-packet
set policy-options community VPN-A members target:65000:1
set routing-instances VRF1 instance-type vrf
set routing-instances VRF1 interface ge-0/0/5.0
set routing-instances VRF1 route-distinguisher 192.168.147.211:1
set routing-instances VRF1 vrf-import VPN-A-import
set routing-instances VRF1 vrf-export VPN-A-export
set routing-instances VRF1 vrf-table-label
set routing-instances VRF1 protocols ospf export bgp-to-ospf
set routing-instances VRF1 protocols ospf area 0.0.0.0 interface ge-0/0/5.0

```

PE2

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.12.2/24

```



```

set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.23.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set protocols mpls static-label-switched-path adj-23 segment 1000123
set protocols mpls static-label-switched-path adj-23 segment next-hop 10.10.23.3
set protocols mpls static-label-switched-path adj-23 segment pop
set protocols mpls static-label-switched-path adj-21 segment 1000221
set protocols mpls static-label-switched-path adj-21 segment next-hop 10.10.12.1
set protocols mpls static-label-switched-path adj-21 segment pop
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0

```

PE3

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.13.3/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.23.3/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.10.34.3/24
set interfaces ge-0/0/2 unit 0 family mpls
set protocols mpls static-label-switched-path adj-34 segment 1000134
set protocols mpls static-label-switched-path adj-34 segment next-hop 10.10.34.4
set protocols mpls static-label-switched-path adj-34 segment pop
set protocols mpls static-label-switched-path adj-32 segment 1000232
set protocols mpls static-label-switched-path adj-32 segment next-hop 10.10.23.2
set protocols mpls static-label-switched-path adj-32 segment pop
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0

```

PE4


```

set interfaces ge-0/0/2 unit 0 family inet address 10.10.34.4/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 10.10.45.4/24
set interfaces ge-0/0/3 unit 0 family mpls
set protocols mpls static-label-switched-path adj-45 segment 1000145
set protocols mpls static-label-switched-path adj-45 segment next-hop 10.10.45.5
set protocols mpls static-label-switched-path adj-45 segment pop
set protocols mpls static-label-switched-path adj-43 segment 1000243
set protocols mpls static-label-switched-path adj-43 segment next-hop 10.10.34.3
set protocols mpls static-label-switched-path adj-43 segment pop
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0

```

PE5

```

set interfaces ge-0/0/3 unit 0 family inet address 10.10.45.5/24
set interfaces ge-0/0/3 unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/4 unit 0 family inet address 10.10.56.5/24
set routing-options autonomous-system 65000
set protocols mpls interface ge-0/0/3.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols bgp group pe type internal
set protocols bgp group pe local-address 192.168.146.181
set protocols bgp group pe family inet-vpn unicast
set protocols bgp group pe neighbor 192.168.147.211
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols bfd sbfd local-discriminator 0.0.0.32 minimum-receive-interval 1000
set protocols source-packet-routing segment-list sl-51 hop-1 ip-address 10.10.45.4
set protocols source-packet-routing segment-list sl-51 hop-2 label 1000243
set protocols source-packet-routing segment-list sl-51 hop-3 label 1000232
set protocols source-packet-routing segment-list sl-51 hop-4 label 1000221
set protocols source-packet-routing source-routing-path lsp-51 to 192.168.147.211
set protocols source-packet-routing source-routing-path lsp-51 primary sl-51
set policy-options policy-statement VPN-A-export term a from protocol ospf
set policy-options policy-statement VPN-A-export term a from protocol direct
set policy-options policy-statement VPN-A-export term a then community add VPN-A

```



```

set policy-options policy-statement VPN-A-export term a then accept
set policy-options policy-statement VPN-A-export term b then reject
set policy-options policy-statement VPN-A-import term a from protocol bgp
set policy-options policy-statement VPN-A-import term a from community VPN-A
set policy-options policy-statement VPN-A-import term a then accept
set policy-options policy-statement VPN-A-import term b then reject
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf from route-filter 10.10.0.0/16 orlonger
set policy-options policy-statement bgp-to-ospf then accept
set policy-options community VPN-A members target:65000:1
set routing-instances VRF1 instance-type vrf
set routing-instances VRF1 interface ge-0/0/4.0
set routing-instances VRF1 route-distinguisher 192.168.146.181:1
set routing-instances VRF1 vrf-import VPN-A-import
set routing-instances VRF1 vrf-export VPN-A-export
set routing-instances VRF1 vrf-table-label
set routing-instances VRF1 protocols ospf export bgp-to-ospf
set routing-instances VRF1 protocols ospf area 0.0.0.0 interface ge-0/0/4.0

```

CE1

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.17.7/24
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0

```

CE2

```

set interfaces ge-0/0/4 unit 0 family inet address 10.10.56.6/24
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0

```

Configuring Device PE1

Step-by-Step Procedure

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

To configure Device PE1:

1. Configure the interfaces.


```
[edit interfaces]
set ge-0/0/0 unit 0 family inet address 10.10.12.1/24
set ge-0/0/0 unit 0 family mpls maximum-labels 5

set ge-0/0/1 unit 0 family inet address 10.10.13.1/24
set ge-0/0/1 unit 0 family mpls maximum-labels 5

set ge-0/0/5 unit 0 family inet address 10.10.17.1/24
```

2. Configure autonomous system number and options to control packet forwarding routing options.

```
[edit routing-options]
set autonomous-system 65000
set forwarding-table export load-balance-policy
set forwarding-table chained-composite-next-hop ingress l3vpn
```

3. Configure the interfaces with the MPLS protocol and configure the MPLS label range.

```
[edit protocols mpls]
set interface ge-0/0/0.0
set interface ge-0/0/1.0
set label-range static-label-range 1000000 1000999
```

4. Configure the type of peer group, local address, protocol family for NLRI in updates, and IP address of a neighbor for the peer group.

```
[edit protocols bgp]
set group pe type internal
set group pe local-address 192.168.147.211
set group pe family inet-vpn unicast
set group pe neighbor 192.168.146.181
```

5. Configure the protocol area interfaces.

```
[edit protocols ospf]
set area 0.0.0.0 interface ge-0/0/0.0
set area 0.0.0.0 interface lo0.0
set area 0.0.0.0 interface ge-0/0/1.0
```


6. Configure the IPv4 address and labels of primary and secondary paths for source routing-traffic engineering (TE) policies of protocol source packet routing (SPRING).

```
[edit protocols source-packet-routing segment-list]
set sl-15-primary hop-1 ip-address 10.10.13.3
set sl-15-primary hop-2 label 1000134
set sl-15-primary hop-3 label 1000145
set sl-15-backup hop-1 ip-address 10.10.12.2
set sl-15-backup hop-2 label 1000123
set sl-15-backup hop-3 label 1000134
set sl-15-backup hop-4 label 1000145
```

7. Configure destination IPv4 address, binding SID label, primary, and secondary source routing path for protocol SPRING.

```
[edit protocols source-packet-routing source-routing-path]
set lsp-15 to 192.168.146.181
set lsp-15 binding-sid 1000999
set lsp-15 primary sl-15-primary
set lsp-15 secondary sl-15-backup
```

8. Configure policy options.

```
[edit policy-options policy-statement]
set VPN-A-export term a from protocol ospf
set VPN-A-export term a from protocol direct
set VPN-A-export term a then community add VPN-A
set VPN-A-export term a then accept
set VPN-A-export term b then reject
set VPN-A-import term a from protocol bgp
set VPN-A-import term a from community VPN-A
set VPN-A-import term a then accept
set VPN-A-import term b then reject
set bgp-to-ospf from protocol bgp
set bgp-to-ospf from route-filter 10.10.0.0/16 orlonger
set bgp-to-ospf then accept
set load-balance-policy then load-balance per-packet
```

9. Configure BGP community information.

```
[edit policy-options]
```



```
set community VPN-A members target:65000:1
```

10. Configure routing instance VRF1 with instance type, interface, router distinguisher, VRF import, export and table label. Configure export policy and interface of area for protocol OSPF.

```
[edit routing-instances VRF1]
set instance-type vrf
set interface ge-0/0/5.0
set route-distinguisher 192.168.147.211:1
set vrf-import VPN-A-import
set vrf-export VPN-A-export
set vrf-table-label
set protocols ospf export bgp-to-ospf
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
```

Results

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

```
user@PE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 55.1.12.1/24;
    }
    family mpls {
      maximum-labels 5;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 55.1.13.1/24;
    }
    family mpls {
      maximum-labels 5;
    }
  }
}
```



```

ge-0/0/5 {
  unit 0 {
    family inet {
      address 55.1.17.1/24;
    }
  }
}

```

user@PE1# show routing-options

```

autonomous-system 65000;
forwarding-table {
  export load-balance-policy;
  chained-composite-next-hop {
    ingress {
      l3vpn;
    }
  }
}

```

user@PE1# show protocols

```

mpls {
  interface ge-0/0/0.0;
  interface ge-0/0/1.0;
  label-range {
    static-label-range 1000000 1000999;
  }
}
bgp {
  group pe {
    type internal;
    local-address 128.9.147.211;
    family inet-vpn {
      unicast;
    }
    neighbor 128.9.146.181;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
    interface ge-0/0/1.0;
  }
}

```



```

    }
}
bfd {
}
source-packet-routing {
    segment-list sl-15-primary {
        hop-1 ip-address 55.1.13.3;
        hop-2 label 1000134;
        hop-3 label 1000145;
    }
    segment-list sl-15-backup {
        hop-1 ip-address 55.1.12.2;
        hop-2 label 1000123;
        hop-3 label 1000134;
        hop-4 label 1000145;
    }
    source-routing-path lsp-15 {
        to 128.9.146.181;
        binding-sid 1000999;
        primary {
            sl-15-primary;
        }
        secondary {
            sl-15-backup;
        }
    }
}
}

```

```

user@PE1# show policy-options
policy-statement VPN-A-export {
    term a {
        from protocol [ ospf direct ];
        then {
            community add VPN-A;
            accept;
        }
    }
    term b {
        then reject;
    }
}
policy-statement VPN-A-import {
    term a {
        from {

```



```

        protocol bgp;
        community VPN-A;
    }
    then accept;
}
term b {
    then reject;
}
}
policy-statement bgp-to-ospf {
    from {
        protocol bgp;
        route-filter 55.1.0.0/16 orlonger;
    }
    then accept;
}
policy-statement load-balance-policy {
    then {
        load-balance per-packet;
    }
}
community VPN-A members target:65000:1;

```

```

user@PE1# show routing-instances
VRF1 {
    instance-type vrf;
    interface ge-0/0/5.0;
    route-distinguisher 128.9.147.211:1;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    vrf-table-label;
    protocols {
        ospf {
            export bgp-to-ospf;
            area 0.0.0.0 {
                interface ge-0/0/5.0;
            }
        }
    }
}

```

Configuring Device PE2

Step-by-Step Procedure

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

1. Configure the interfaces.

```
[edit interfaces]
set ge-0/0/0 unit 0 family inet address 10.10.12.2/24
set ge-0/0/0 unit 0 family mpls

set ge-0/0/1 unit 0 family inet address 10.10.23.2/24
set ge-0/0/1 unit 0 family mpls
```

2. Configure the static LSP for protocol MPLS.

```
[edit protocols mpls static-label-switched-path]
set adj-23 segment 1000123
set adj-23 segment next-hop 10.10.23.3
set adj-23 segment pop
set adj-21 segment 1000221
set adj-21 segment next-hop 10.10.12.1
set adj-21 segment pop
```

3. Configure interfaces and static label range for protocol MPLS.

```
[edit protocols mpls]
set interface ge-0/0/0.0
set interface ge-0/0/1.0
set label-range static-label-range 1000000 1000999
```

4. Configure interfaces for protocol OSPF.

```
[edit protocols ospf area 0.0.0.0]
set interface ge-0/0/0.0
set interface ge-0/0/1.0
```

Results

From configuration mode on router PE2, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

user@PE2# **show interfaces**

```
ge-0/0/0 {
  unit 0 {
    family inet {
      address 55.1.12.2/24;
    }
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 55.1.23.2/24;
    }
    family mpls;
  }
}
```

user@PE2# **show protocols**

```
mpls {
  static-label-switched-path adj-23 {
    segment {
      1000123;
      next-hop 55.1.23.3;
      pop;
    }
  }
  static-label-switched-path adj-21 {
    segment {
      1000221;
      next-hop 55.1.12.1;
      pop;
    }
  }
  interface ge-0/0/0.0;
  interface ge-0/0/1.0;
  label-range {
    static-label-range 1000000 1000999;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
```



```
}
}
```

Verification

IN THIS SECTION

- Verifying Route Entry of Routing Table inet.3 of Router PE1 | 671
- Verifying Route Table Entries of Routing Table mpls.0 of Router PE1 | 672
- Verifying SPRING Traffic Engineered LSP of Router PE1 | 672
- Verifying SPRING Traffic Engineered LSPs on the Ingress Router of Router PE1 | 673
- Verifying the Routing Table Entries of Routing Table mpls.0 of Router PE2 | 674
- Verifying the Status of Static MPLS LSP Segments of Router PE2 | 675

Confirm that the configuration is working properly.

Verifying Route Entry of Routing Table inet.3 of Router PE1

Purpose

Verify the route entry of routing table inet.3 of router PE1.

Action

From operational mode, enter the **show route table inet.3** command.

```
user@PE1> show route table inet.3
```

```
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.146.181/32    *[SPRING-TE/8] 03:09:26, metric 1
                    > to 10.10.13.3 via ge-0/0/1.0, Push 1000145, Push 1000134(top)

                    to 10.10.12.2 via ge-0/0/0.0, Push 1000145, Push 1000134,
                    Push 1000123(top)
```

Meaning

The output displays the ingress routes of segment routing tunnels.

Verifying Route Table Entries of Routing Table mpls.0 of Router PE1

Purpose

Verify the route entries of routing table mpls.0

Action

From operational mode, enter the **show route table mpls.0** command.

user@PE1> **show route table mpls.0**

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 03:25:52, metric 1
           Receive
1          *[MPLS/0] 03:25:52, metric 1
           Receive
2          *[MPLS/0] 03:25:52, metric 1
           Receive
13         *[MPLS/0] 03:25:52, metric 1
           Receive
16         *[VPN/0] 03:25:52
           > via lsi.0 (VRF1), Pop
1000999    *[SPRING-TE/8] 03:04:03, metric 1
           > to 10.10.13.3 via ge-0/0/1.0, Swap 1000145, Push 1000134(top)

           to 10.10.12.2 via ge-0/0/0.0, Swap 1000145, Push 1000134,
           Push 1000123(top)
```

Meaning

The output displays the SID labels of segment routing tunnels.

Verifying SPRING Traffic Engineered LSP of Router PE1

Purpose

Verify SPRING traffic engineered LSPs on the ingress routers.

Action

From operational mode, enter the **show spring-traffic-engineering overview** command.

user@PE1> **show spring-traffic-engineering overview**


```

Overview of SPRING-TE:
  Route preference: 8
  Number of LSPs: 1 (Up: 1, Down: 0)
  External controllers:
    < Not configured >

```

Meaning

The output displays the overview of SPRING traffic engineered LSPs on the ingress router.

Verifying SPRING Traffic Engineered LSPs on the Ingress Router of Router PE1

Purpose

Verify SPRING traffic engineered LSPs on the ingress router.

Action

From operational mode, enter the **show spring-traffic-engineering lsp detail** command.

```
user@PE1# show spring-traffic-engineering lsp detail
```

```

Name: lsp-15
To: 192.168.146.181
State: Up
  Path: sl-15-primary
  Outgoing interface: ge-0/0/1.0
  BFD status: N/A (Up: 0, Down: 0)
  SR-ERO hop count: 3
    Hop 1 (Strict):
      NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.10.13.3
      SID type: None
    Hop 2 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 1000134
    Hop 3 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 1000145
  Path: sl-15-backup
  Outgoing interface: ge-0/0/0.0
  BFD status: N/A (Up: 0, Down: 0)
  SR-ERO hop count: 4
    Hop 1 (Strict):
      NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.10.12.2
      SID type: None

```



```

Hop 2 (Strict):
  NAI: None
  SID type: 20-bit label, Value: 1000123
Hop 3 (Strict):
  NAI: None
  SID type: 20-bit label, Value: 1000134
Hop 4 (Strict):
  NAI: None
  SID type: 20-bit label, Value: 1000145

Total displayed LSPs: 1 (Up: 1, Down: 0)

```

Meaning

The output displays details of SPRING traffic engineered LSPs on the ingress router

Verifying the Routing Table Entries of Routing Table mpls.0 of Router PE2

Purpose

Verify the routing table entries of routing table mpls.0 of router PE2.

Action

From operational mode, enter the **show route table mpls.0** command.

```
user@PE2> 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] 03:22:29, metric 1
           Receive
1          *[MPLS/0] 03:22:29, metric 1
           Receive
2          *[MPLS/0] 03:22:29, metric 1
           Receive
13         *[MPLS/0] 03:22:29, metric 1
           Receive
1000123    *[MPLS/6] 03:22:29, metric 1
           > to 10.10.23.3 via ge-0/0/1.0, Pop
1000123(S=0) *[MPLS/6] 03:22:29, metric 1
           > to 10.10.23.3 via ge-0/0/1.0, Pop
1000221    *[MPLS/6] 03:22:29, metric 1
           > to 10.10.12.1 via ge-0/0/0.0, Pop

```



```
1000221(S=0)          *[MPLS/6] 03:22:29, metric 1
                      > to 10.10.12.1 via ge-0/0/0.0, Pop
```

Verifying the Status of Static MPLS LSP Segments of Router PE2

Purpose

Verify the status of MPLS LSP segments of router PE2.

Action

From operational mode, enter the **show mpls static-lsp** command.

user@PE2> **show mpls static-lsp**

```
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0

Segment LSPs:
LSPname                SID-label    State
adj-21                 1000221      Up
adj-23                 1000123      Up
Total 2, displayed 2, Up 2, Down 0
```

Meaning

The output displays the status of static MPLS LSP segments of router PE2.

Release History Table

Release	Description
19.4R1	You can configure a tunnel template for PCE-initiated segment routing LSPs to pass down two additional parameters for these LSPs - Bidirectional forwarding detection (BFD) and LDP tunneling.
19.1R1	Starting in Junos OS Release 19.1R1, a commit check feature is introduced to ensure that all the segment lists contributing to the colored routes have the minimum label present for all hops.
19.1R1	Starting in Junos OS Release 19.1R1, this requirement does not apply, as the first hop of the non-colored static LSPs now provides support for SID labels, in addition to IP addresses. With the first hop label support, MPLS fast reroute (FRR) and weighted equal-cost multipath is enabled for resolving the static non-colored segment routing LSPs, similar to colored static LSPs.
18.2R1	Starting in Junos OS Release 18.2R1, statically configured non-colored segment routing LSPs on the ingress device are reported to the Path Computation Element (PCE) through a Path Computation Element Protocol (PCEP) session.

RELATED DOCUMENTATION

[segment](#) | [2768](#)

[segment-list](#) | [2770](#)

[source-routing-path](#) | [2782](#)

Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution

IN THIS SECTION

- [Understanding RE-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution](#) | [677](#)
- [Configuring RE-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution](#) | [679](#)
- [Example](#) | [681](#)
- [Verify That LSPs Are Configured for Static Segment-Routing Tunnels and That S-BFD Session Status Is Visible](#) | [683](#)

- [Verify the Segment-Routing Tunnel Route with a Primary Next Hop and a Secondary Next Hop | 684](#)
- [Verify the S-BFD Session of the Primary Path | 685](#)

You can run seamless Bidirectional Forwarding Detection (S-BFD) over non-colored and colored label-switched paths (LSPs) with first-hop label resolution, using S-BFD as a fast mechanism to detect path failures.

Understanding RE-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution

IN THIS SECTION

- [S-BFD Static Segment-Routing Tunnels for First-Hop Labels | 677](#)
- [Limitations | 679](#)

S-BFD Static Segment-Routing Tunnels for First-Hop Labels

IN THIS SECTION

- [LSP-level S-BFD | 678](#)
- [S-BFD Parameters | 678](#)

Segment-routing architecture enables ingress nodes in a core network to steer traffic through explicit paths through the network. The segment-routing traffic engineering (TE) next hop is a list or lists of segment identifiers (SIDs). These segment lists represent paths in the network that you want incoming traffic to take. The incoming traffic may be labeled or IP traffic and the forwarding operation at the ingress node may be a label swap or a destination-based lookup to steer the traffic onto these segment-routing TE paths in the forwarding path.

You can run seamless BFD (S-BFD) over non-colored and colored static segment-routing LSPs with first-hop label resolution and use S-BFD as a fast mechanism to detect path failures and to trigger global convergence. S-BFD requires fewer state changes than BFD requires, thus speeding up the reporting of path failures.

Given a segment-routing tunnel with one or multiple primary LSPs and optionally a secondary LSP, you can enable S-BFD on any of those LSPs. If an S-BFD session goes down, the software updates the segment-routing tunnel's route by deleting the next hops of the failed LSPs. If the first-hop label of the LSP points to more than one immediate next hop, the kernel continues to send S-BFD packets if *at least* one next hop is available (underlying next-hop reachability failure detection must be faster than the duration of the S-BFD detection timer).

NOTE: This model is supported for auto-translate-derived LSPs.

LSP-level S-BFD

An S-BFD session is created for each unique label-stack+address-family combination. You can configure identical segment lists and enable S-BFD for all of them. The segment lists that have identical label-stack+address-family combinations share the same S-BFD session. The source address for the S-BFD session is set to the least configured loopback address (except the special addresses) under the main instance.

NOTE: Ensure that the chosen source address is routable.

The address family of an LSP is derived based on the address family of the “to” address in the segment-routing TE tunnel. The state of the LSP with S-BFD configured also depends on whether BFD is up—if S-BFD is configured for an LSP, the LSP route isn't calculated until S-BFD reports the path is alive.

S-BFD Parameters

The following S-BFD parameters are supported for segment-routing TE paths:

- Remote discriminator
- Minimum interval
- Multiplier
- No router alert option

In S-BFD, each responder may have multiple discriminators. The discriminators may be advertised by IGP to other routers, or they may be statically configured on these routers. On an initiator, a particular discriminator is chosen as the remote discriminator for an S-BFD session by static configuration, based on the decision or resolution made by you or by a central controller. When multiple LSPs are created with the same key label stack and S-BFD is enabled on each of them with different parameters, the aggressive value of each parameter is used in the shared S-BFD session. For the discriminator parameter, the lowest value is considered as most aggressive. Similarly for the router alert option, if one of the configurations no router alert is configured, the derived S-BFD parameter will have no router alert option.

Limitations

- Only global repair is supported.
- Even though S-BFD detects failures depending on the configured timer values, convergence time depends on the global repair time (*seconds*).

Configuring RE-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution

To enable LSP-level S-BFD for a segment list, you configure the **bfd-liveness-detection** configuration statement at the **[edit protocols source-packet-routing source-routing-path *lsp-path-name* primary *segment-list-name*]** hierarchy and the **[edit protocols source-packet-routing source-routing-path *lsp-path-name* secondary *segment-list-name*]** hierarchy.

The following steps give the basic configuration and then operation of S-BFD with first-hop label resolution:

- The steps immediately below describe the outlines of the basic *configuration*:
 1. On an ingress router, you configure one or more segment lists with first-hop labels for a static segment-routing tunnel to use as primary and secondary paths.
 2. On the ingress router, you configure the static segment-routing tunnel, with either multiple primary paths (for load balancing), or one primary path and one secondary path (for path protection). Each primary or secondary path (LSP) refers to one of the segment lists you configured already, creating routes using the next hops derived from the first-hop labels from contributing LSPs.
 3. On the ingress router, you enable per-packet load-balancing so that routes resolving over ingress routes and the binding-SID route (which all have first-hop labels) install all active paths in the Packet Forwarding Engine. An S-BFD session under an LSP protects all routes that use that LSP.
 4. On the egress router of the segment-routing tunnel, you configure S-BFD responder mode with a local discriminator D, creating a distributed S-BFD listener session for D on each FPC.
 5. On the ingress router, you configure S-BFD for any LSP for which you want to see path-failure detection. You specify a remote-discriminator D to match the local S-BFD discriminator of the egress router. An S-BFD initiator session is created with the LSP label-stack+address-family combination as the key, if an initiator session doesn't already exist for the current LSP key. The S-BFD parameters in the case of a matching BFD session are reevaluated with the new shared LSPs taken into consideration. When the S-BFD parameters are derived, the aggressive value of each parameter is chosen.

The steps immediately below describe basic *operation* :

1. The S-BFD initiator session runs in a centralized mode on the Routing Engine. The software tracks S-BFD session up and down states.
2. When the S-BFD state moves to UP, the LSP is considered for the relevant routes.

3. On the ingress router, if the software detects an S-BFD session DOWN, a session-down notification is sent to the routing daemon (RPD), and RPD deletes the next hop of the failed LSPs from the segment-routing tunnel's route.
4. The total traffic loss in the procedure is bound to the S-BFD failure detection time and the global repair time. The S-BFD failure detection time is determined by the S-BFD minimum-interval and multiplier parameters. The global repair time depends on the segment-routing TE process time and the redownload of the routes to forwarding.

LSP label stacks are changed as follows:

1. The particular LSP is detached from the existing S-BFD session. If the existing S-BFD session has at least one LSP referring to it, the old BFD session is preserved, but the S-BFD parameters are re-evaluated so that the aggressive parameter values among the existing LSP sessions is chosen. Also, the name of the S-BFD session is updated to the least one if there is a change. If the old S-BFD session has no more LSPs attached to it, that S-BFD session is removed.
2. The software attempts to find an existing BFD session that matches the new-label-stack+address-family combination value; if such a match exists, the LSP refers to that existing S-BFD session. The S-BFD session is re-evaluated for any change in parameters or session name similarly to the re-evaluations in step 1.
3. If there is no matching BFD session in the system, a new BFD session is created, and the S-BFD parameters are derived from this LSP.

NOTE: An S-BFD session's minimum interval and multiplier determine the failure detection time for the session. The repair time additionally depends on the global repair time.

The following output shows configuration statements you would use for a colored LSP with primary LSPs:

```
[edit protocols]
source-packet-routing {
  source-routing-path lsp_name {
    to destination_address;
    color color_value;
    binding-sid binding-label;
    primary segment_list_1_name weight weight;
    ... {
      bfd-liveness-detection {
        sbfd {
          remote-discriminator value;
        }
      }
    }
  }
}
```



```

    }
}

```

At the responder side, you must configure the correct discriminator:

```

[edit protocols bfd]
sbfd {
    local-discriminator value;
}

```

By default, router alerts are configured for S-BFD packets. When the MPLS header is removed at the responder end, the packet is sent to the host for processing without a destination address lookup for the packet. If you enable the no-router-alert option on the ingress router, the router-alert option is removed from the IP options and hence from the egress side. You must configure the destination address explicitly in lo0; otherwise the packet is discarded, and S-BFD remains down.

```

[edit interfaces lo0 unit 0 family inet]
address 127.0.0.1/32;

```

You can use a new trace flag, **bfd**, to trace BFD activities:

```

user@host# set protocols source-packet-routing traceoptions flag bfd

```

Example

The following configuration is an example of a non-colored static segment-routing tunnel with LSP protection.

```

protocols {
    source-packet-routing {
        source-routing-path ncsrlsp5 {
            to 10.10.10.10;
            primary {
                ncsrpath12 {
                    weight 1;
                    bfd-liveness-detection {
                        sbfd {
                            remote-discriminator 100;
                        }
                    }
                    minimum-interval 100;
                }
            }
        }
    }
}

```



```

}
ncsrpath13 {
  weight 2;
  bfd-liveness-detection {
    sbfd {
      remote-discriminator 100;
    }
    minimum-interval 100;
  }
}
ncsrpath14 {
  weight 3;
  bfd-liveness-detection {
    sbfd {
      remote-discriminator 100;
    }
    minimum-interval 100;
  }
}
ncsrpath15 {
  weight 4;
  bfd-liveness-detection {
    sbfd {
      remote-discriminator 100;
    }
    minimum-interval 100;
  }
}
segment-list ncsrpath12 {
  hop1 label 50191;
  hop2 label 801000;
}
segment-list ncsrpath13 {
  hop1 label 50191;
  hop2 label 801001;
  hop3 label 801000;
}
segment-list ncsrpath14 {
  hop1 label 801000;
}
segment-list ncsrpath15 {
  hop1 label 801002;
  hop2 label 801000;
}

```



```

    }
  }
}
}

```

Verify That LSPs Are Configured for Static Segment-Routing Tunnels and That S-BFD Session Status Is Visible

Purpose

Use the **show spring-traffic-engineering lsp detail** command to show LSPs for static segment-routing tunnels, with S-BFD session status.

Action

```
user@host> show spring-traffic-engineering lsp detail
```

```

Name: abc
To: 77.77.77.77
State: Up
  Path: s11
  Outgoing interface: NA
  BFD status: Up BFD name: V4-s11
  SR-ERO hop count: 3
    Hop 1 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 801007
    Hop 2 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 22222
    Hop 3 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 3333
  Path: s12
  Outgoing interface: NA
  BFD status: Up BFD name: V4-s12
  SR-ERO hop count: 2
    Hop 1 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 801006
    Hop 2 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 121212
  Path: s12

```



```

Outgoing interface: NA
BFD status: Up BFD name: V4-sl2
SR-ERO hop count: 2
  Hop 1 (Strict):
    NAI: None
    SID type: 20-bit label, Value: 801006
  Hop 2 (Strict):
    NAI: None
    SID type: 20-bit label, Value: 121212

Total displayed LSPs: 1 (Up: 1, Down: 0)

```

Because many LSPs can share the same BFD session, when the first LSP with a unique label-stack+address-family combination comes up, the name of the S-BFD session uses address-family+lsp-name. If more LSPs later share the same session, the name of the session can change to address-family+least-lsp-name, without affecting the state of the S-BFD session. The name of the S-BFD session appears in output from the **show bfd session extensive** command as well. Output for each LSP shows the S-BFD status as well as the S-BFD name it is referring to as shown in the preceding example as **BFD status: Up BFD name: V4-sl2**. Because there might not be one S-BFD session per LSP, the LSP-level S-BFD counters are not displayed.

Verify the Segment-Routing Tunnel Route with a Primary Next Hop and a Secondary Next Hop

Purpose

On the Routing Engine of the ingress router, verify the segment-routing tunnel route with a primary next hop and a secondary next hop.

Action

```

user@host> show route table inet.3

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

128.9.146.157/32    *[SPRING-TE/8] 00:43:16, metric 1
                   > to 55.1.12.2 via ge-0/0/0.0, Push 1000145, Push 1000134,
Push 1000123(top)
                   to 55.1.12.2 via ge-0/0/1.0, Push 1000934, Push 1000923(top)

```


Verify the S-BFD Session of the Primary Path

Purpose

On the Routing Engine of the ingress router, verify the S-BFD session of the primary path.

Action

```
user@host> show bfd session extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
127.0.0.1	Up		4.000	1.000	4

```

Client SPRING-TE, TX interval 1.000, RX interval 1.000
Session up time 00:40:53, previous down time 00:02:08
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Session type: Multi hop BFD (Seamless)
Min async interval 1.000, min slow interval 1.000
Adaptive async TX interval 1.000, RX interval 1.000
Local min TX interval 1.000, minimum RX interval 1.000, multiplier 4
Remote min TX interval 1.000, min RX interval 0.001, multiplier 4
Local discriminator 28, remote discriminator 32
Echo mode disabled/inactive
Remote is control-plane independent
Path-Name V4-sl-1

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

```

NOTE: On the Routing Engine of the ingress router, verify the S-BFD session of the secondary path also similarly.

RELATED DOCUMENTATION

[bfd-liveness-detection \(LSP\) | 2571](#)

[sbfd | 2764](#)

Configuring Explicit-Path LSPs

If you disable constrained-path label-switched path (LSP) computation, as described in [“Disabling Constrained-Path LSP Computation” on page 489](#), you can configure LSPs manually or allow the LSPs to follow the IGP path.

When explicit-path LSPs are configured, the LSP is established along the path you specified. If the path is topologically not feasible, either because the network is partitioned or insufficient resources are available along some parts of the path, the LSP will fail. No alternative paths can be used. If the setup succeeds, the LSP stays on the defined path indefinitely.

To configure an explicit-path LSP, follow these steps:

1. Configure the path information in a named path, as described in [“Creating Named Paths” on page 516](#). To configure complete path information, specify every router hop between the ingress and egress routers, preferably using the **strict** attribute. To configure incomplete path information, specify only a subset of router hops, using the **loose** attribute in places where the path is incomplete.

For incomplete paths, the MPLS routers complete the path by querying the local routing table. This query is done on a hop-by-hop basis, and each router can figure out only enough information to reach the next explicit hop. It might be necessary to traverse a number of routers to reach the next (loose) explicit hop.

Configuring incomplete path information creates portions of the path that depend on the current routing table, and this portion of the path can reroute itself as the topology changes. Therefore, an explicit-path LSP that contains incomplete path information is not completely fixed. These types of LSPs have only a limited ability to repair themselves, and they tend to create loops or flaps depending on the contents of the local routing table.

2. To configure the LSP and point it to the named path, use either the **primary** or **secondary** statement, as described in [“Configuring Primary and Secondary LSPs” on page 565](#).
3. Disable constrained-path LSP computation by including the **no-cspf** statement either as part of the LSP or as part of a **primary** or **secondary** statement. For more information, see [“Disabling Constrained-Path LSP Computation” on page 489](#).
4. Configure any other LSP properties.

Using explicit-path LSPs has the following drawbacks:

- More configuration effort is required.
- Configured path information cannot take into account dynamic network bandwidth reservation, so the LSPs tend to fail when resources become depleted.
- When an explicit-path LSP fails, you might need to manually repair it.

Because of these limitations, we recommend that you use explicit-path LSPs only in controlled situations, such as to enforce an optimized LSP placement strategy resulting from computations with an offline simulation software package.

Example: Configuring an Explicit-Path LSP

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

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



```

    }
    interface so-0/0/0;
  }
}

```

Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP

In a network where aggregate Ethernet (AE) bundles are in use, an aggregate link could be bundle of any number of physical links. The traffic sent over these AE bundle interfaces are forwarded on any of the member links of an AE interface. The traffic can take any physical link based on the hash defined for load-balancing the traffic, which makes it difficult to isolate which links have gone bad or are dropping the traffic. One way to test the forwarding path available in the network is to add a single-hop static label switched path (LSP) with the next hop pointing to a specific member link of the AE interface.

The default label operation for the static LSPs must be pop and forward. When a packet hits these labeled routes, the packet is forwarded on to a specific member link. A unique label is used to identify the member link. These labels are referred to as adjacency segment identifiers (SID) and are statically provisioned.

You can control the flow of the packets in the network by constructing a label stack in controller, which includes the labels allocated by all transit static LSP. Operation, Administration, and Maintenance (OAM) packets are crafted and injected into the network with entire label-stack.

When a packet hits this label route the label is popped and traffic is forwarded on the member link specified in the configuration. A destination MAC is chosen while forwarding the packet, the destination Mac is the aggregate interface MAC address (selected based on nexthop address configured).

When the member link goes down and aggregate interface is up, then the route corresponding to that member link is deleted. When an aggregate interface goes down, then all the routes corresponding to member links of the aggregate interface are deleted. When the child physical interface is LACP detached but the child physical interface is up, the labeled route for the child link is deleted. In the case of LACP detach, if the member link is up and invalid forwarding state, then the OAM packets is dropped in the PFE when the child physical interface is detached.

Use the following example to configure single-hop static LSP for an AE member link.

1. Define a static label range.

```
user@host# set protocols mpls label-range static-label-range 1000000 1048575;
```


NOTE: We recommend configuring the default static label range of 1000000-1048575 for the static LSP. If you wish to configure a label range other than the default static label range, configure multiple ranges.

2. Create a static LSP for the AE member link from the segment routing local block (SRLB) pool of the static label range.

```
user@host# set protocols mpls static-label-switched-path static-lsp transit 100001 pop next-hop 10.1.1.1 member-interface ge-0/0/0
```

In this configuration, a transit labelled router is installed in mpls.0, pops the label, and forwards the packet down the next hop. The next hop address is mandatory for broadcast interfaces (such as ge-, xe-, ae-) and the if-name is used for P2P interfaces (such as so-). The address is required for broadcast interfaces because the next hop IP address is used to pick the destination MAC address. The source MAC address for the packet is the AE's MAC address.

The sample outputs display the member link name in the next hop output:

show mpls static-lsp extensive

```
user@host> show mpls static-lsp extensive
```

```
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0

Transit LSPs:
LSPname: static-lsp1, Incoming-label: 1000001
  Description: verify-static-lsp-behavior
  State: Up, Sub State: Traffic via primary but unprotected
  Nexthop: 10.2.1.1 Via ae0.0 -> ge-0/0/0
  LabelOperation: Pop
  Created: Thu May 25 15:31:26 2017
  Bandwidth: 0 bps
  Statistics: Packets 0, Bytes 0
```

show route label label-name extensive

```
user@host> show route label 1000001 extensive
```

```
mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
```



```

1000001 (1 entry, 1 announced)
TSI:
KRT in-kernel 1000001/52 -> {Pop      }
    *MPLS   Preference: 6
            Next hop type: Router, Next hop index: 611
            Address: 0xb7a17b0
            Next-hop reference count: 2
            Next hop: 10.2.1.1 via ae0.0 -> ge-0/0/0 weight 0x1, selected
            Label operation: Pop
            Load balance label: None;
            Label element ptr: 0xb7a1800
            Label parent element ptr: 0x0
            Label element references: 1
            Label element child references: 0
            Label element lsp id: 0
            Session Id: 0x15d
            State: <Active Int>
            Age: 3:13:15      Metric: 1
            Validation State: unverified
            Task: MPLS
            Announcement bits (1): 1-KRT
            AS path: I
            Label 188765184

```

RELATED DOCUMENTATION

| [transit](#) | 2898

Configuring Point-to-Multipoint LSPs

IN THIS CHAPTER

- [Point-to-Multipoint LSPs Overview | 691](#)
- [Understanding Point-to-Multipoint LSPs | 693](#)
- [Point-to-Multipoint LSP Configuration Overview | 694](#)
- [Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP | 695](#)
- [Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs | 720](#)
- [Configuring Inter-Domain Point-to-Multipoint LSPs | 722](#)
- [Configuring Link Protection for Point-to-Multipoint LSPs | 724](#)
- [Configuring Graceful Restart for Point-to-Multipoint LSPs | 724](#)
- [Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs | 725](#)
- [Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs | 726](#)
- [Enabling Point-to-Point LSPs to Monitor Egress PE Routers | 727](#)
- [Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases | 728](#)


Point-to-Multipoint LSPs Overview

A point-to-multipoint MPLS LSP is an LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in [Figure 51 on page 692](#). Router PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Router PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Router P1 replicates the packet and forwards it to Routers PE2 and PE3. Router P2 sends the packet to Router PE4.

This feature is described in detail in the Internet drafts [draft-raggarwa-mpls-p2mp-te-02.txt](#) (expired February 2004), *Establishing Point to Multipoint MPLS TE LSPs*, [draft-ietf-mpls-rsvp-te-p2mp-02.txt](#), *Extensions to Resource Reservation Protocol-Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label-Switched Paths (LSPs)*, and RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths* (only point-to-multipoint LSPs are supported).

Figure 51: Point-to-Multipoint LSPs



The following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP enables you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- You can configure a node to be both a transit and an egress router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any of the primary paths fail, traffic can be quickly switched to the bypass.
- You can configure branch LSPs either statically, dynamically, or as a combination of static and dynamic LSPs.
- You can enable graceful Routing Engine switchover (GRES) and graceful restart for point-to-multipoint LSPs at ingress and egress routers. The point-to-multipoint LSPs must be configured using either static routes or circuit cross-connect (CCC). GRES and graceful restart allow the traffic to be forwarded at the Packet Forwarding Engine based on the old state while the control plane recovers. Feature parity for GRES and graceful restart for MPLS point-to-multipoint LSPs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

RELATED DOCUMENTATION

High Availability User Guide


Example: NG-VPLS Using Point-to-Multipoint LSPs

Understanding Point-to-Multipoint LSPs

A point-to-multipoint MPLS label-switched path (LSP) is an LDP-signaled or RSVP-signaled LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the inbound (ingress) router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in [Figure 52 on page 693](#). Device PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Device PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Device P1 replicates the packet and forwards it to Routers PE2 and PE3. Device P2 sends the packet to Device PE4.

Figure 52: Point-to-Multipoint LSPs



Following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP allows you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- You can configure a node to be both a transit and an outbound (egress) router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any primary paths fail, traffic can be quickly switched to the bypass.
- You can configure subpaths either statically or dynamically.
- You can enable graceful restart on point-to-multipoint LSPs.

RELATED DOCUMENTATION

[MPLS Traffic Engineering and Signaling Protocols Overview | 1134](#)

[Point-to-Multipoint LSP Configuration Overview | 694](#)

Point-to-Multipoint LSP Configuration Overview

To set up a point-to-multipoint LSP:

1. Configure the primary LSP from the ingress router and the branch LSPs that carry traffic to the egress routers.
2. Specify a pathname on the primary LSP and this same path name on each branch LSP.

NOTE: By default, the branch LSPs are dynamically signaled by means of Constrained Shortest Path First (CSPF) and require no configuration. You can alternatively configure the branch LSPs as static paths.

RELATED DOCUMENTATION

Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP

IN THIS SECTION

- [Requirements | 695](#)
- [Overview | 695](#)
- [Configuration | 696](#)
- [Verification | 718](#)

This example shows how to configure a collection of paths to create an RSVP-signaled point-to-multipoint label-switched path (LSP).

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, multiple routing devices serve as the transit, branch, and leaf nodes of a single point-to-multipoint LSP. On the provider edge (PE), Device PE1 is the ingress node. The branches go from PE1 to PE2, PE1 to PE3, and PE1 to PE4. Static unicast routes on the ingress node (PE1) point to the egress nodes.


This example also demonstrates static routes with a next hop that is a point-to-multipoint LSP, using the [p2mp-lsp-next-hop](#) statement. This is useful when implementing filter-based forwarding.

NOTE: Another option is to use the **lsp-next-hop** statement to configure a regular point-to-point LSP to be the next hop. Though not shown in this example, you can optionally assign an independent preference and metric to the next hop.

Topology Diagram

Figure 53 on page 696 shows the topology used in this example.

Figure 53: RSVP-Signaled Point-to-Multipoint LSP



Configuration

CLI Quick Configuration

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

Device PE1

```
set interfaces ge-2/0/2 unit 0 description PE1-to-CE1
set interfaces ge-2/0/2 unit 0 family inet address 10.0.244.10/30
set interfaces fe-2/0/10 unit 1 description PE1-to-P2
set interfaces fe-2/0/10 unit 1 family inet address 2.2.2.1/24
set interfaces fe-2/0/10 unit 1 family mpls
set interfaces fe-2/0/9 unit 8 description PE1-to-P3
set interfaces fe-2/0/9 unit 8 family inet address 6.6.6.1/24
set interfaces fe-2/0/9 unit 8 family mpls
set interfaces fe-2/0/8 unit 9 description PE1-to-P4
set interfaces fe-2/0/8 unit 9 family inet address 3.3.3.1/24
```



```

set interfaces fe-2/0/8 unit 9 family mpls
set interfaces lo0 unit 1 family inet address 100.10.10.10/32
set protocols rsvp interface fe-2/0/10.1
set protocols rsvp interface fe-2/0/9.8
set protocols rsvp interface fe-2/0/8.9
set protocols rsvp interface lo0.1
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path PE1-PE2 to 100.50.50.50
set protocols mpls label-switched-path PE1-PE2 link-protection
set protocols mpls label-switched-path PE1-PE2 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE3 to 100.70.70.70
set protocols mpls label-switched-path PE1-PE3 link-protection
set protocols mpls label-switched-path PE1-PE3 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE4 to 100.40.40.40
set protocols mpls label-switched-path PE1-PE4 link-protection
set protocols mpls label-switched-path PE1-PE4 p2mp p2mp1
set protocols mpls interface fe-2/0/10.1
set protocols mpls interface fe-2/0/9.8
set protocols mpls interface fe-2/0/8.9
set protocols mpls interface lo0.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface fe-2/0/10.1
set protocols ospf area 0.0.0.0 interface fe-2/0/9.8
set protocols ospf area 0.0.0.0 interface fe-2/0/8.9
set protocols ospf area 0.0.0.0 interface lo0.1
set routing-options static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
set routing-options router-id 100.10.10.10

```

Device CE1

```

set interfaces ge-1/3/2 unit 0 family inet address 10.0.244.9/30
set interfaces ge-1/3/2 unit 0 description CE1-to-PE1
set routing-options static route 10.0.104.8/30 next-hop 10.0.244.10
set routing-options static route 10.0.134.8/30 next-hop 10.0.244.10
set routing-options static route 10.0.224.8/30 next-hop 10.0.244.10

```

Device CE2


```

set interfaces ge-1/3/3 unit 0 family inet address 10.0.224.9/30
set interfaces ge-1/3/3 unit 0 description CE2-to-PE2
set routing-options static route 10.0.244.8/30 next-hop 10.0.224.10

```

Device CE3

```

set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.9/30
set interfaces ge-2/0/1 unit 0 description CE3-to-PE3
set routing-options static route 10.0.244.8/30 next-hop 10.0.134.10

```

Device CE4

```

set interfaces ge-3/1/3 unit 0 family inet address 10.0.104.10/30
set interfaces ge-3/1/3 unit 0 description CE4-to-PE4
set routing-options static route 10.0.244.8/30 next-hop 10.0.104.9

```

Configuring the Ingress Label-Switched Router (LSR) (Device PE1)

Step-by-Step Procedure

To configure Device PE1:

1. Configure the interfaces, interface encapsulation, and protocol families.

```

[edit interfaces]
user@PE1# set ge-2/0/2 unit 0 description PE1-to-CE1
user@PE1# set ge-2/0/2 unit 0 family inet address 10.0.244.10/30
user@PE1# set fe-2/0/10 unit 1 description PE1-to-P2
user@PE1# set fe-2/0/10 unit 1 family inet address 2.2.2.1/24
user@PE1# set fe-2/0/10 unit 1 family mpls
user@PE1# set fe-2/0/9 unit 8 description PE1-to-P3
user@PE1# set fe-2/0/9 unit 8 family inet address 6.6.6.1/24
user@PE1# set fe-2/0/9 unit 8 family mpls
user@PE1# set fe-2/0/8 unit 9 description PE1-to-P4
user@PE1# set fe-2/0/8 unit 9 family inet address 3.3.3.1/24
user@PE1# set fe-2/0/8 unit 9 family mpls
user@PE1# set lo0 unit 1 family inet address 100.10.10.10/32

```


2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit protocols]
user@PE1# set rsvp interface fe-2/0/10.1
user@PE1# set rsvp interface fe-2/0/9.8
user@PE1# set rsvp interface fe-2/0/8.9
user@PE1# set rsvp interface lo0.1
user@PE1# set mpls interface fe-2/0/10.1
user@PE1# set mpls interface fe-2/0/9.8
user@PE1# set mpls interface fe-2/0/8.9
user@PE1# set mpls interface lo0.1
user@PE1# set ospf area 0.0.0.0 interface ge-2/0/2.0
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/10.1
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/9.8
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/8.9
user@PE1# set ospf area 0.0.0.0 interface lo0.1
```

3. Configure the MPLS point-to-multipoint LSPs.

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-PE2 to 100.50.50.50
user@PE1# set mpls label-switched-path PE1-PE2 p2mp p2mp1
user@PE1# set mpls label-switched-path PE1-PE3 to 100.70.70.70
user@PE1# set mpls label-switched-path PE1-PE3 p2mp p2mp1
user@PE1# set mpls label-switched-path PE1-PE4 to 100.40.40.40
user@PE1# set mpls label-switched-path PE1-PE4 p2mp p2mp1
```

4. (Optional) Enable link protection on the LSPs.

Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-PE2 link-protection
user@PE1# set mpls label-switched-path PE1-PE3 link-protection
user@PE1# set mpls label-switched-path PE1-PE4 link-protection
```

5. Enable MPLS to perform traffic engineering for OSPF.

```
[edit protocols]
user@PE1# set mpls traffic-engineering bgp-igp
```


This causes the ingress routes to be installed in the inet.0 routing table. By default, MPLS performs traffic engineering for BGP only. You need to enable MPLS traffic engineering on the ingress LSR only.

6. Enable traffic engineering for OSPF.

```
[edit protocols]
user@PE1# set ospf traffic-engineering
```

This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

7. Configure the router ID.

```
[edit routing-options]
user@PE1# set router-id 100.10.10.10
```

8. Configure static IP unicast routes with the point-to-multipoint LSP name as the next hop for each route.

```
[edit routing-options]
user@PE1# set static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
user@PE1# set static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
user@PE1# set static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
```

9. If you are done configuring the device, commit the configuration.

```
[edit]
user@PE1# commit
```

Configuring the Transit and Egress LSRs (Devices P2, P3, P4, PE2, PE3, and PE4)

Step-by-Step Procedure

To configure the transit and egress LSRs:

1. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit]
user@P2# set interfaces fe-2/0/10 unit 2 description P2-to-PE1
user@P2# set interfaces fe-2/0/10 unit 2 family inet address 2.2.2.2/24
user@P2# set interfaces fe-2/0/10 unit 2 family mpls
user@P2# set interfaces fe-2/0/9 unit 10 description P2-to-PE2
```



```

user@P2# set interfaces fe-2/0/9 unit 10 family inet address 5.5.5.1/24
user@P2# set interfaces fe-2/0/9 unit 10 family mpls
user@P2# set interfaces lo0 unit 2 family inet address 100.20.20.20/32
user@PE2# set interfaces ge-2/0/3 unit 0 description PE2-to-CE2
user@PE2# set interfaces ge-2/0/3 unit 0 family inet address 10.0.224.10/30
user@PE2# set interfaces fe-2/0/10 unit 5 description PE2-to-P2
user@PE2# set interfaces fe-2/0/10 unit 5 family inet address 5.5.5.2/24
user@PE2# set interfaces fe-2/0/10 unit 5 family mpls
user@PE2# set interfaces lo0 unit 5 family inet address 100.50.50.50/32
user@P3# set interfaces fe-2/0/10 unit 6 description P3-to-PE1
user@P3# set interfaces fe-2/0/10 unit 6 family inet address 6.6.6.2/24
user@P3# set interfaces fe-2/0/10 unit 6 family mpls
user@P3# set interfaces fe-2/0/9 unit 11 description P3-to-PE3
user@P3# set interfaces fe-2/0/9 unit 11 family inet address 7.7.7.1/24
user@P3# set interfaces fe-2/0/9 unit 11 family mpls
user@P3# set interfaces lo0 unit 6 family inet address 100.60.60.60/32
user@PE3# set interfaces ge-2/0/1 unit 0 description PE3-to-CE3
user@PE3# set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.10/30
user@PE3# set interfaces fe-2/0/10 unit 7 description PE3-to-P3
user@PE3# set interfaces fe-2/0/10 unit 7 family inet address 7.7.7.2/24
user@PE3# set interfaces fe-2/0/10 unit 7 family mpls
user@PE3# set interfaces lo0 unit 7 family inet address 100.70.70.70/32
user@P4# set interfaces fe-2/0/10 unit 3 description P4-to-PE1
user@P4# set interfaces fe-2/0/10 unit 3 family inet address 3.3.3.2/24
user@P4# set interfaces fe-2/0/10 unit 3 family mpls
user@P4# set interfaces fe-2/0/9 unit 12 description P4-to-PE4
user@P4# set interfaces fe-2/0/9 unit 12 family inet address 4.4.4.1/24
user@P4# set interfaces fe-2/0/9 unit 12 family mpls
user@P4# set interfaces lo0 unit 3 family inet address 100.30.30.30/32
user@PE4# set interfaces ge-2/0/0 unit 0 description PE4-to-CE4
user@PE4# set interfaces ge-2/0/0 unit 0 family inet address 10.0.104.9/30
user@PE4# set interfaces fe-2/0/10 unit 4 description PE4-to-P4
user@PE4# set interfaces fe-2/0/10 unit 4 family inet address 4.4.4.2/24
user@PE4# set interfaces fe-2/0/10 unit 4 family mpls
user@PE4# set interfaces lo0 unit 4 family inet address 100.40.40.40/32

```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

```

[edit]
user@P2# set protocols rsvp interface fe-2/0/10.2
user@P2# set protocols rsvp interface fe-2/0/9.10
user@P2# set protocols rsvp interface lo0.2
user@P2# set protocols mpls interface fe-2/0/10.2

```



```

user@P2# set protocols mpls interface fe-2/0/9.10
user@P2# set protocols mpls interface lo0.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/9.10
user@P2# set protocols ospf area 0.0.0.0 interface lo0.2
user@PE2# set protocols rsvp interface fe-2/0/10.5
user@PE2# set protocols rsvp interface lo0.5
user@PE2# set protocols mpls interface fe-2/0/10.5
user@PE2# set protocols mpls interface lo0.5
user@PE2# set protocols ospf area 0.0.0.0 interface ge-2/0/3.0
user@PE2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.5
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.5
user@P3# set protocols rsvp interface fe-2/0/10.6
user@P3# set protocols rsvp interface fe-2/0/9.11
user@P3# set protocols rsvp interface lo0.6
user@P3# set protocols mpls interface fe-2/0/10.6
user@P3# set protocols mpls interface fe-2/0/9.11
user@P3# set protocols mpls interface lo0.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/9.11
user@P3# set protocols ospf area 0.0.0.0 interface lo0.6
user@PE3# set protocols rsvp interface fe-2/0/10.7
user@PE3# set protocols rsvp interface lo0.7
user@PE3# set protocols mpls interface fe-2/0/10.7
user@PE3# set protocols mpls interface lo0.7
user@PE3# set protocols ospf area 0.0.0.0 interface ge-2/0/1.0
user@PE3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.7
user@PE3# set protocols ospf area 0.0.0.0 interface lo0.7
user@P4# set protocols rsvp interface fe-2/0/10.3
user@P4# set protocols rsvp interface fe-2/0/9.12
user@P4# set protocols rsvp interface lo0.3
user@P4# set protocols mpls interface fe-2/0/10.3
user@P4# set protocols mpls interface fe-2/0/9.12
user@P4# set protocols mpls interface lo0.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/9.12
user@P4# set protocols ospf area 0.0.0.0 interface lo0.3
user@PE4# set protocols rsvp interface fe-2/0/10.4
user@PE4# set protocols rsvp interface lo0.4
user@PE4# set protocols mpls interface fe-2/0/10.4
user@PE4# set protocols mpls interface lo0.4
user@PE4# set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
user@PE4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.4
user@PE4# set protocols ospf area 0.0.0.0 interface lo0.4

```


3. Enable traffic engineering for OSPF.

```
[edit]
user@P2# set protocols ospf traffic-engineering
user@P3# set protocols ospf traffic-engineering
user@P4# set protocols ospf traffic-engineering
user@PE2# set protocols ospf traffic-engineering
user@PE3# set protocols ospf traffic-engineering
user@PE4# set protocols ospf traffic-engineering
```

This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

4. Configure the router IDs.

```
[edit]
user@P2# set routing-options router-id 100.20.20.20
user@P3# set routing-options router-id 100.60.60.60
user@P4# set routing-options router-id 100.30.30.30
user@PE2# set routing-options router-id 100.50.50.50
user@PE3# set routing-options router-id 100.70.70.70
user@PE4# set routing-options router-id 100.40.40.40
```

5. If you are done configuring the devices, commit the configuration.

```
[edit]
user@host# commit
```

Results

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

Device PE1

```
user@PE1# show interfaces
ge-2/0/2 {
  unit 0 {
    description R1-to-CE1;
```



```

        family inet {
            address 10.0.244.10/30;
        }
    }
}
fe-2/0/10 {
    unit 1 {
        description PE1-to-P2;
        family inet {
            address 2.2.2.1/24;
        }
        family mpls;
    }
}
fe-2/0/9 {
    unit 8 {
        description PE1-to-P2;
        family inet {
            address 6.6.6.1/24;
        }
        family mpls;
    }
}
fe-2/0/8 {
    unit 9 {
        description PE1-to-P3;
        family inet {
            address 3.3.3.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 1 {
        family inet {
            address 100.10.10.10/32;
        }
    }
}

```

user@PE1# **show protocols**


```

rsvp {
  interface fe-2/0/10.1;
  interface fe-2/0/9.8;
  interface fe-2/0/8.9;
  interface lo0.1;
}
mpls {
  traffic-engineering bgp-igp;
  label-switched-path PE1-to-PE2 {
    to 100.50.50.50;
    link-protection;
    p2mp p2mp1;
  }
  label-switched-path PE1-to-PE3 {
    to 100.70.70.70;
    link-protection;
    p2mp p2mp1;
  }
  label-switched-path PE1-to-PE4 {
    to 100.40.40.40;
    link-protection;
    p2mp p2mp1;
  }
  interface fe-2/0/10.1;
  interface fe-2/0/9.8;
  interface fe-2/0/8.9;
  interface lo0.1;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-2/0/2.0;
    interface fe-2/0/10.1;
    interface fe-2/0/9.8;
    interface fe-2/0/8.9;
    interface lo0.1;
  }
}
}

```

```

user@PE1# show routing-options
static {
  route 5.5.5.0/24 {
    p2mp-lsp-next-hop p2mp1;
  }
}

```



```

route 7.7.7.0/24 {
    p2mp-lsp-next-hop p2mp1;
}
route 4.4.4.0/24 {
    p2mp-lsp-next-hop p2mp1;
}
}
router-id 100.10.10.10;

```

Device P2

```

user@P2# show interfaces
fe-2/0/10 {
    unit 2 {
        description P2-to-PE1;
        family inet {
            address 2.2.2.2/24;
        }
        family mpls;
    }
}
fe-2/0/9 {
    unit 10 {
        description P2-to-PE2;
        family inet {
            address 5.5.5.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 2 {
        family inet {
            address 100.20.20.20/32;
        }
    }
}
}

```

```

user@P2# show protocols
rsvp {
    interface fe-2/0/10.2;
}

```



```

    interface fe-2/0/9.10;
    interface lo0.2;
}
mpls {
    interface fe-2/0/10.2;
    interface fe-2/0/9.10;
    interface lo0.2;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-2/0/10.2;
        interface fe-2/0/9.10;
        interface lo0.2;
    }
}
}

```

```

user@P2# show routing-options
router-id 100.20.20.20;

```

Device P3

```

user@P3# show interfaces
fe-2/0/10 {
    unit 6 {
        description P3-to-PE1;
        family inet {
            address 6.6.6.2/24;
        }
        family mpls;
    }
}
fe-2/0/9 {
    unit 11 {
        description P3-to-PE3;
        family inet {
            address 7.7.7.1/24;
        }
        family mpls;
    }
}
}

```



```

lo0 {
  unit 6 {
    family inet {
      address 100.60.60.60/32;
    }
  }
}

```

user@P3# **show protocols**

```

rsvp {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
mpls {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fe-2/0/10.6;
    interface fe-2/0/9.11;
    interface lo0.6;
  }
}

```

user@P2# **show routing-options**

```

router-id 100.60.60.60;

```

Device P4

user@P4# **show interfaces**

```

fe-2/0/10 {
  unit 3 {
    description P4-to-PE1;
    family inet {

```



```

        address 3.3.3.2/24;
    }
    family mpls;
}
fe-2/0/9 {
    unit 12 {
        description P4-to-PE4;
        family inet {
            address 4.4.4.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 3 {
        family inet {
            address 100.30.30.30/32;
        }
    }
}
}

```

user@P4# **show protocols**

```

rsvp {
    interface fe-2/0/10.3;
    interface fe-2/0/9.12;
    interface lo0.3;
}
mpls {
    interface fe-2/0/10.3;
    interface fe-2/0/9.12;
    interface lo0.3;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-2/0/10.3;
        interface fe-2/0/9.12;
        interface lo0.3;
    }
}
}

```



```
user@P3# show routing-options
router-id 100.30.30.30;
```

Device PE2

```
user@PE2# show interfaces
ge-2/0/3 {
  unit 0 {
    description PE2-to-CE2;
    family inet {
      address 10.0.224.10/30;
    }
  }
}
fe-2/0/10 {
  unit 5 {
    description PE2-to-P2;
    family inet {
      address 5.5.5.2/24;
    }
    family mpls;
  }
}
lo0 {
  unit 5 {
    family inet {
      address 100.50.50.50/32;
    }
  }
}
```

```
user@PE2# show protocols
rsvp {
  interface fe-2/0/10.5;
  interface lo0.5;
}
mpls {
  interface fe-2/0/10.5;
  interface lo0.5;
```



```

}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-2/0/3.0;
    interface fe-2/0/10.5;
    interface lo0.5;
  }
}

```

```

user@PE2# show routing-options
router-id 100.50.50.50;

```

Device PE3

```

user@PE3# show interfaces
ge-2/0/1 {
  unit 0 {
    description PE3-to-CE3;
    family inet {
      address 10.0.134.10/30;
    }
  }
}
fe-2/0/10 {
  unit 7 {
    description PE3-to-P3;
    family inet {
      address 7.7.7.2/24;
    }
    family mpls;
  }
}
lo0 {
  unit 7 {
    family inet {
      address 100.70.70.70/32;
    }
  }
}
}

```



```
user@PE3# show protocols
```

```

rsvp {
  interface fe-2/0/10.7;
  interface lo0.7;
}
mpls {
  interface fe-2/0/10.7;
  interface lo0.7;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-2/0/1.0;
    interface fe-2/0/10.7;
    interface lo0.7;
  }
}

```

```
user@PE3# show routing-options
```

```
router-id 100.70.70.70;
```

Device PE4

```
user@PE4# show interfaces
```

```

ge-2/0/0 {
  unit 0 {
    description PE4-to-CE4;
    family inet {
      address 10.0.104.9/30;
    }
  }
}
fe-2/0/10 {
  unit 4 {
    description PE4-to-P4;
    family inet {
      address 4.4.4.2/24;
    }
    family mpls;
  }
}

```



```

    lo0 {
      unit 4 {
        family inet {
          address 100.40.40.40/32;
        }
      }
    }
  }
}

```

```
user@PE4# show protocols
```

```

rsvp {
  interface fe-2/0/10.4;
  interface lo0.4;
}
mpls {
  interface fe-2/0/10.4;
  interface lo0.4;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-2/0/0.0;
    interface fe-2/0/10.4;
    interface lo0.4;
  }
}

```

```
user@PE4# show routing-options
```

```
router-id 100.40.40.40;
```

Configuring Device CE1

Step-by-Step Procedure

To configure Device CE1:

1. Configure an interface to Device PE1.

```

[edit interfaces]
user@CE1# set ge-1/3/2 unit 0 family inet address 10.0.244.9/30
user@CE1# set ge-1/3/2 unit 0 description CE1-to-PE1

```


2. Configure static routes from Device CE1 to the three other customer networks, with Device PE1 as the next hop.

```
[edit routing-options]
user@CE1# set static route 10.0.104.8/30 next-hop 10.0.244.10
user@CE1# set static route 10.0.134.8/30 next-hop 10.0.244.10
user@CE1# set static route 10.0.224.8/30 next-hop 10.0.244.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE1# commit
```

Results

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

```
user@CE1# show interfaces
ge-1/3/2 {
  unit 0 {
    family inet {
      address 10.0.244.9/30;
      description CE1-to-PE1;
    }
  }
}
```

```
user@CE1# show routing-options
static {
  route 10.0.104.8/30 next-hop 10.0.244.10;
  route 10.0.134.8/30 next-hop 10.0.244.10;
  route 10.0.224.8/30 next-hop 10.0.244.10;
}
```

Configuring Device CE2

Step-by-Step Procedure

To configure Device CE2:

1. Configure an interface to Device PE2.

```
[edit interfaces]
user@CE2# set ge-1/3/3 unit 0 family inet address 10.0.224.9/30
user@CE2# set ge-1/3/3 unit 0 description CE2-to-PE2
```

2. Configure a static route from Device CE2 to CE1, with Device PE2 as the next hop.

```
[edit routing-options]
user@CE2# set static route 10.0.244.8/30 next-hop 10.0.224.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE2# commit
```

Results

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

```
user@CE2# show interfaces
ge-1/3/3 {
  unit 0 {
    family inet {
      address 10.0.224.9/30;
      description CE2-to-PE2;
    }
  }
}
```

```
user@CE2# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.224.10;
}
```


Configuring Device CE3

Step-by-Step Procedure

To configure Device CE3:

1. Configure an interface to Device PE3.

```
[edit interfaces]
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.134.9/30
user@CE3# set ge-2/0/1 unit 0 description CE3-to-PE3
```

2. Configure a static route from Device CE3 to CE1, with Device PE3 as the next hop.

```
[edit routing-options]
user@CE3# set static route 10.0.244.8/30 next-hop 10.0.134.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE3# commit
```

Results

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

```
user@CE3# show interfaces
ge-2/0/1 {
  unit 0 {
    family inet {
      address 10.0.134.9/30;
      description CE3-to-PE3;
    }
  }
}
```

```
user@CE3# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.134.10;
}
```


Configuring Device CE4

Step-by-Step Procedure

To configure Device CE4:

1. Configure an interface to Device PE4.

```
[edit interfaces]
user@CE4# set ge-3/1/3 unit 0 family inet address 10.0.104.10/30
user@CE4# set ge-3/1/3 unit 0 description CE4-to-PE4
```

2. Configure a static route from Device CE4 to CE1, with Device PE4 as the next hop.

```
[edit routing-options]
user@CE4# set static route 10.0.244.8/30 next-hop 10.0.104.9
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE4# commit
```

Results

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

```
user@CE4# show interfaces
ge-3/1/3 {
  unit 0 {
    family inet {
      address 10.0.104.10/30;
      description CE4-to-PE4;
    }
  }
}
```

```
user@CE4# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.104.9;
}
```


Verification

IN THIS SECTION

- [Verifying Connectivity | 718](#)
- [Verifying the State of the Point-to-Multipoint LSP | 719](#)
- [Checking the Forwarding Table | 719](#)

Confirm that the configuration is working properly.

Verifying Connectivity

Purpose

Make sure that the devices can ping each other.

Action

Run the **ping** command from CE1 to the interface on CE2 connecting to PE2.

```
user@CE1> ping 10.0.224.9
```

```
PING 10.0.224.9 (10.0.224.9): 56 data bytes
64 bytes from 10.0.224.9: icmp_seq=0 ttl=61 time=1.387 ms
64 bytes from 10.0.224.9: icmp_seq=1 ttl=61 time=1.394 ms
64 bytes from 10.0.224.9: icmp_seq=2 ttl=61 time=1.506 ms
^C
--- 10.0.224.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.387/1.429/1.506/0.055 ms
```

Run the **ping** command from CE1 to the interface on CE3 connecting to PE3.

```
user@CE1> ping 10.0.134.9
```

```
PING 10.0.134.9 (10.0.134.9): 56 data bytes
64 bytes from 10.0.134.9: icmp_seq=0 ttl=61 time=1.068 ms
64 bytes from 10.0.134.9: icmp_seq=1 ttl=61 time=1.062 ms
64 bytes from 10.0.134.9: icmp_seq=2 ttl=61 time=1.053 ms
^C
--- 10.0.134.9 ping statistics ---
```



```
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.053/1.061/1.068/0.006 ms
```

Run the **ping** command from CE1 to the interface on CE4 connecting to PE4.

```
user@CE1> ping 10.0.104.10
```

```
PING 10.0.104.10 (10.0.104.10): 56 data bytes
64 bytes from 10.0.104.10: icmp_seq=0 ttl=61 time=1.079 ms
64 bytes from 10.0.104.10: icmp_seq=1 ttl=61 time=1.048 ms
64 bytes from 10.0.104.10: icmp_seq=2 ttl=61 time=1.070 ms
^C
--- 10.0.104.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.048/1.066/1.079/0.013 ms
```

Verifying the State of the Point-to-Multipoint LSP

Purpose

Make sure that the ingress, transit, and egress LSRs are in the Up state.

Action

Run the **show mpls lsp p2mp** command on all of the LSRs. Only the ingress LSR is shown here.

```
user@PE1> show mpls lsp p2mp
```

```
Ingress LSP: 1 sessions
P2MP name: p2mp1, P2MP branch count: 3
To          From          State Rt P    ActivePath    LSPname
100.40.40.40 100.10.10.10 Up      0 *          PE1-PE4
100.70.70.70 100.10.10.10 Up      0 *          PE1-PE3
100.50.50.50 100.10.10.10 Up      0 *          PE1-PE2
Total 3 displayed, Up 3, Down 0
...
```

Checking the Forwarding Table

Purpose

Make sure that the routes are set up as expected by running the **show route forwarding-table** command. Only the routes to the remote customer networks are shown here.

Action

user@PE1> show route forwarding-table

```
Routing table: default.inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
...
10.0.104.8/30         user    0 3.3.3.2          ucst  1006    6 fe-2/0/8.9
10.0.134.8/30         user    0 6.6.6.2          ucst  1010    6 fe-2/0/9.8
10.0.224.8/30         user    0 2.2.2.2          ucst  1008    6 fe-2/0/10.1
...
```

RELATED DOCUMENTATION

| *MPLS Applications User Guide*

Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs

IN THIS SECTION

- [Configuring the Primary Point-to-Multipoint LSP | 721](#)
- [Configuring a Branch LSP for Point-to-Multipoint LSPs | 721](#)

A point-to-multipoint MPLS label-switched path (LSP) is an RSVP LSP with multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. For more information about point-to-multipoint LSPs, see [“Point-to-Multipoint LSPs Overview” on page 691](#).

To configure a point-to-multipoint LSP, you need to configure the primary LSP from the ingress router and the branch LSPs that carry traffic to the egress routers, as described in the following sections:

Configuring the Primary Point-to-Multipoint LSP

A point-to-multipoint LSP must have a configured primary point-to-multipoint LSP to carry traffic from the ingress router. The configuration of the primary point-to-multipoint LSP is similar to a signaled LSP. See [“Configuring the Ingress Router for MPLS-Signaled LSPs” on page 516](#) for more information. In addition to the conventional LSP configuration, you need to specify a path name for the primary point-to-multipoint LSP by including the **p2mp** statement:

```
p2mp p2mp-lsp-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

You can enable the optimization timer for point-to-multipoint LSPs. See [“Optimizing Signaled LSPs” on page 541](#) for more information.

Configuring a Branch LSP for Point-to-Multipoint LSPs

IN THIS SECTION

- [Configuring the Branch LSP as a Dynamic Path | 722](#)
- [Configuring the Branch LSP as a Static Path | 722](#)

The primary point-to-multipoint LSP sends traffic to two or more branch LSPs carrying traffic to each of the egress provider edge (PE) routers. In the configuration for each of these branch LSPs, the point-to-multipoint LSP path name you specify must be identical to the path name configured for the primary point-to-multipoint LSP. See [“Configuring the Primary Point-to-Multipoint LSP” on page 721](#) for more information.

To associate a branch LSP with the primary point-to-multipoint LSP, specify the point-to-multipoint LSP name by including the **p2mp** statement:

```
p2mp p2mp-lsp-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]

- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]

NOTE: Any change in any of the branch LSPs of a point-to-multipoint LSP, either due to a user action or an automatic adjustment made by the router, causes the primary and branch LSPs to be resigned. The new point-to-multipoint LSP is signaled first before the old path is taken down.

The following sections describe how you can configure the branch LSP as a dynamically signaled path using Constrained Shortest Path First (CSPF), as a static path, or as a combination of dynamic and static paths:

Configuring the Branch LSP as a Dynamic Path

By default, the branch LSP for a point-to-multipoint LSP is signaled dynamically using CSPF and requires no configuration.

When a point-to-multipoint LSP is changed, either by the addition or deletion of new destinations or by the recalculation of the path to existing destinations, certain nodes in the tree might receive data from more than one incoming interface. This can happen under the following conditions:

- Some of the branch LSPs to destinations are statically configured and might intersect with statically or dynamically calculated paths to other destinations.
- When a dynamically calculated path for a branch LSP results in a change of incoming interface for one of the nodes in the network, the older path is not immediately torn down after the new one has been signaled. This ensures that any data in transit relying on the older path can reach its destination. However, network traffic can potentially use either path to reach the destination.
- A faulty router at the ingress calculates the paths to two different branch destinations such that a different incoming interface is chosen for these branch LSPs on a router node common to these branch LSPs.

Configuring the Branch LSP as a Static Path

You can configure the branch LSP for a point-to-multipoint LSP as a static path. See [“Configuring Static LSPs” on page 603](#) for more information.

Configuring Inter-Domain Point-to-Multipoint LSPs

An inter-domain P2MP LSP is a P2MP LSP that has one or more sub-LSPs (branches) that span multiple domains in a network. Examples of such domains include IGP areas and autonomous systems (ASs). A sub-LSP of an inter-domain P2MP LSP may be intra-area, inter-area, or inter-AS, depending on the location of the egress node (leaf) with respect to the ingress node (source).

On the ingress node, a name is assigned to the inter-domain P2MP LSP and shared by all constituent sub-LSPs. Each sub-LSP is configured separately, with its own egress node and optionally an explicit path. The location of the egress node of the sub-LSP with respect to the ingress node determines whether the sub-LSP is intra-area, inter-area, or inter-AS.

Inter-domain P2MP LSPs can be used to transport traffic in the following applications in a multi-area or multi-AS network:

- Layer 2 broadcast and multicast over MPLS
- Layer 3 BGP/MPLS VPN
- VPLS

On each domain boundary node (ABR or ASBR) along the path of the P2MP LSP, the **expand-loose-hop** statement must be configured at the **[edit protocols mpls]** hierarchy level so that CSPF can extend a loose-hop ERO (usually the first entry of the ERO list carried by RSVP Path message) towards the egress node or the next domain boundary node.

CSPF path computation for inter-domain P2MP LSPs:

- CSPF path computation is supported on each sub-LSP for inter-domain P2MP LSPs. A sub-LSP may be intra-area, inter-area, or inter-AS. CSPF treats an inter-area or inter-AS sub-LSP in the same manner as an inter-domain P2P LSP.
- On an ingress node or a domain boundary node (ABR or ASBR), CSPF can perform an Explicit Route Object (ERO) expansion per-RSVP query. The destination queried could be an egress node or a received loose-hop ERO. If the destination resides in a neighboring domain that the node is connected to, CSPF generates either a sequence of strict-hop EROs towards it or a sequence of strict-hop EROs towards another domain boundary node that can reach the destination.
- If RSVP fails to signal a path through a previously selected domain boundary node, RSVP attempts to signal a path through other available domain boundary nodes in a round-robin fashion.
- When a sub-LSP is added or removed to or from an inter-domain P2MP LSP, causing its path (branch) to be merged or pruned with or from the current P2MP tree, the paths being taken by the other sub-LSPs should not be affected, helping to prevent traffic disruption on those sub-LSPs.

Be aware of the following when deploying inter-domain P2MP LSPs in your network:

- Periodic path re-optimization is supported for inter-domain P2MP LSPs on ingress nodes. It can be turned on for an inter-domain P2MP LSP by configuring the **optimize-timer** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level with the same interval for every sub-LSP.
- Only link protection bypass LSPs are supported for inter-domain P2MP LSPs. To enable it for an inter-domain P2MP LSP, link-protection must be configured for all sub-LSPs and on all of the RSVP interfaces that the P2MP LSP might travel through.
- Only OSPF areas are supported for inter-domain P2MP LSPs. IS-IS levels are not supported.

Configuring Link Protection for Point-to-Multipoint LSPs

Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. When link protection is configured for an interface and a point-to-multipoint LSP that traverses this interface, a bypass LSP is created that handles this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination.

To extend link protection to all of the paths used by a point-to-multipoint LSP, link protection must be configured on each router that each branch LSP traverses. If you enable link protection on a point-to-multipoint LSP, you must enable link protection on all of the branch LSPs.

The Internet draft [draft-ietf-mpls-rsvp-te-p2mp-01.txt](#), *Extensions to RSVP-TE for Point to Multipoint TE LSPs*, describes link protection for point-to-multipoint LSPs.

To enable link protection on point-to-multipoint LSPs, complete the following steps:

1. Configure link protection on each branch LSP. To configure link protection, include the **link-protection** statement:

```
link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *branch-lsp-name*]
 - [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *branch-lsp-name*]
2. Configure link protection for each RSVP interface on each router that the branch LSP traverses. For information about how to configure link protection on RSVP interfaces, see [“Configuring Link Protection on Interfaces Used by LSPs” on page 871](#).

For more information on how to configure link protection, see [“Configuring Node Protection or Link Protection for LSPs” on page 879](#).

Configuring Graceful Restart for Point-to-Multipoint LSPs

You can configure graceful restart on point-to-multipoint LSPs. Graceful restart allows a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not apparent to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers.

To enable graceful restart on a router handling point-to-multipoint LSP traffic, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

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

The graceful restart configuration for point-to-multipoint LSPs is identical to that of point-to-point LSPs. For more information on how to configure graceful restart, see [“Configuring RSVP Graceful Restart” on page 883](#).

Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs

You can control whether a reverse path forwarding (RPF) check is performed for a source and group entry before installing a route in the multicast forwarding cache. This makes it possible to use point-to-multipoint LSPs to distribute multicast traffic to PIM islands situated downstream from the egress routers of the point-to-multipoint LSPs.

By configuring the **rpf-check-policy** statement, you can disable RPF checks for a source and group pair. You would typically configure this statement on the egress routers of a point-to-multipoint LSP, because the interface receiving the multicast traffic on a point-to-multipoint LSP egress router might not always be the RPF interface.

You can also configure a routing policy to act upon a source and group pair. This policy behaves like an import policy, so if no policy term matches the input data, the default policy action is “acceptance.” An accept policy action enables RPF checks. A reject policy action (applied to all source and group pairs that are not accepted) disables RPF checks for the pair.

To configure a multicast RPF check policy for a point-to-multipoint LSP, specify the RPF check policy using the **rpf-check-policy** statement:

```
rpf-check-policy policy;
```

You can include this statement at the following hierarchy levels:

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

You also must configure a policy for the multicast RPF check. You configure policies at the **[edit policy-options]** hierarchy level. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

NOTE: When you configure the **rpf-check-policy** statement, the Junos OS cannot perform RPF checks on incoming traffic and therefore cannot detect traffic arriving on the wrong interface. This might cause routing loops to form.

Example: Configuring Multicast RPF Check Policy for a Point-to-Multipoint LSP

Configure a policy to ensure that an RPF check is not performed for sources with prefix **128.83/16** or longer that belong to groups having a prefix of **228/8** or longer:

```
[edit]
policy-options {
  policy-statement rpf-sg-policy {
    from {
      route-filter 228.0.0.0/8 orlonger;
      source-address-filter 128.83.0.0/16 orlonger;
    }
    then {
      reject;
    }
  }
}
```

Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs

You can configure one or more PE routers as part of a backup PE router group to enable ingress PE router redundancy. You accomplish this by configuring the IP addresses of the backup PE routers (at least one backup PE router is required) and the local IP address used by the local PE router.

You must also configure a full mesh of point-to-point LSPs between the primary and backup PE routers. You also need to configure BFD on these LSPs. See [“Configuring BFD for RSVP-Signaled LSPs” on page 110](#) and [“Configuring BFD for LDP LSPs” on page 975](#) for more information.

To configure ingress PE router redundancy for point-to-multipoint LSPs, include the **backup-pe-group** statement:

```
backup-pe-group pe-group-name {
    backups [addresses];
    local-address address;
}
```

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

After you configure the ingress PE router redundancy backup group, you must also apply the group to a static route on the PE router. This ensures that the static route is active (installed in the forwarding table) when the local PE router is the designated forwarder for the backup PE group. You can only associate a backup PE router group with a static route that also has the **p2mp-lsp-next-hop** statement configured. For more information, see [“Configuring Static Unicast Routes for Point-to-Multipoint LSPs” on page 611](#).

Enabling Point-to-Point LSPs to Monitor Egress PE Routers

Configuring an LSP with the **associate-backup-pe-groups** statement enables it to monitor the status of the PE router to which it is configured. You can configure multiple backup PE router groups using the same router's address. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. The **associate-backup-pe-groups** statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to that address.

To allow an LSP to monitor the status of the egress PE router, include the **associate-backup-pe-groups** statement:

```
associate-backup-pe-groups;
```

This statement can be configured at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

If you configure the **associate-backup-pe-groups** statement, you must configure BFD for the point-to-point LSP. For information about how to configure BFD for an LSP, see [“Configuring BFD for MPLS IPv4 LSPs” on page 109](#) and [“Configuring BFD for LDP LSPs” on page 975](#).

You also must configure a full mesh of point-to-point LSPs between the PE routers in the backup PE router group. A full mesh is required so that each PE router within the group can independently determine the

status of the other PE routers, allowing each router to independently determine which PE router is currently the designated forwarder for the backup PE router group.

If you configure multiple LSPs with the **associate-backup-pe-groups** statement to the same destination PE router, the first LSP configured is used to monitor the forwarding state to that PE router. If you configure multiple LSPs to the same destination, make sure to configure similar parameters for the LSPs. With this configuration scenario, a failure notification might be triggered even though the remote PE router is still up.

Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases

In Junos OS Release 9.1 and earlier, Resv messages that include the S2L_SUB_LSP object are rejected by default. In Junos OS Release 9.2 and later, such messages are accepted by default. To ensure proper functioning of point-to-multipoint LSPs in a network that includes both devices running Junos OS Release 9.1 and earlier and devices running Junos 9.2 and later, you must include the **no-p2mp-sublsp** statement in the configuration of the devices running Junos 9.2 and later:

```
no-p2mp-sublsp;
```

You can include this statement at the following hierarchy levels:

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

Configuring Container LSPs

IN THIS CHAPTER

- [Dynamic Bandwidth Management Using Container LSP Overview | 729](#)
- [Example: Configuring Dynamic Bandwidth Management Using Container LSP | 759](#)
- [Configuring Dynamic Bandwidth Management Using Container LSP | 791](#)

Dynamic Bandwidth Management Using Container LSP Overview

IN THIS SECTION

- [Understanding RSVP Multipath Extensions | 730](#)
- [Junos OS RSVP Multipath Implementation | 731](#)
- [Current Traffic Engineering Challenges | 731](#)
- [Using Container LSP as a Solution | 734](#)
- [Junos OS Container LSP Implementation | 736](#)
- [Configuration Statements Supported for Container LSPs | 753](#)
- [Impact of Configuring Container LSPs on Network Performance | 757](#)
- [Supported and Unsupported Features | 758](#)

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

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.

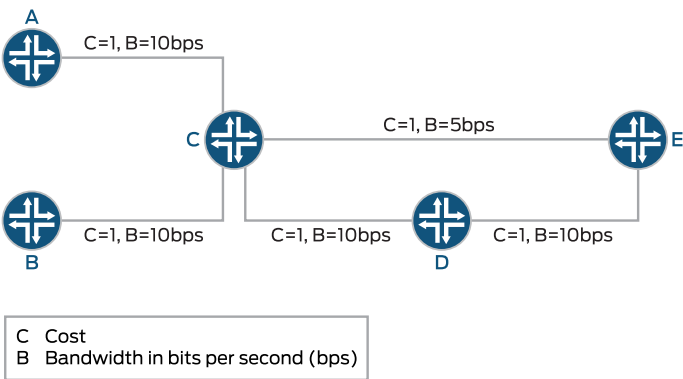
Junos OS implements the RSVP multipath extensions without the need for protocol extensions. A single container LSP, which has the characteristics of ECMP and RSVP TE, is provisioned. A container LSP consists of several member LSPs and is set up between the ingress and egress routing device. Each member 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 form available resources.

Figure 54 on page 731 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 54: Sample Topology



Some of the traffic engineering problems seen in [Figure 54 on page 731](#) 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 17 on page 732](#).

Table 17: 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 54 on page 731](#), 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 54 on page 731](#), the LSPs follow the sequence mentioned in [Table 18 on page 732](#).

Table 18: 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 19 on page 734](#) and [Table 20 on page 734](#). The ERO that an LSP uses depends on its signaling time.

Table 19: 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 20: 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

IN THIS SECTION

- [Accommodating the New Demand X | 735](#)
- [Creating New LSPs to Meet Demand X | 735](#)
- [Assigning Bandwidth to the New LSPs | 735](#)
- [Controlling the LSP Paths | 735](#)

A container LSP can be used as a solution to the challenges faced by the current traffic engineering features. Considering [Figure 54 on page 731](#), 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

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

IN THIS SECTION

- [Container LSP Terminology | 736](#)
- [LSP Splitting | 737](#)
- [LSP Merging | 740](#)
- [Node and Link Protection | 741](#)
- [Naming Convention | 742](#)
- [Normalization | 742](#)
- [Constraint-Based Routing Path Computation | 748](#)
- [Sampling | 749](#)
- [Support for NSR, IPG-FA, and Static Routes | 749](#)

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

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

IN THIS SECTION

- [Operational Overview | 737](#)
- [Operational Constraints | 738](#)
- [Supported Criteria | 739](#)
- [Splitting Triggers | 739](#)

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-signaled with different bandwidths due to the autobandwidth adjustments. If a container LSP is not configured with autobandwidth, then the member LSPs are signaled with the static bandwidth value, if configured. There is no dynamic splitting in this case, as there is no dynamic estimation of aggregate bandwidth. The splitting adjustments with a specific bandwidth value can be manually triggered.

NOTE:

Be aware of the following considerations for LSP splitting:

- After LSP splitting, the ingress router continues to inject one forwarding adjacency. Forwarding adjacencies are not supported in IGP for this feature.
- Between two normalization events, two LSPs might have different bandwidths subjected to autobandwidth constraints.
- After LSPs are split (or merged), make-before-break uses the fixed filter (FF) style sharing unless the **adaptive** option is configured. However, two different LSPs do not do the shared explicit (SE) style sharing for this feature.
- When LSPs are re-signaled with modified bandwidths, some of the LSPs might not get signaled successfully, leading to failover options.

Operational Constraints

LSP splitting has the following operational constraints:

- LSP bandwidth—Although there are a number of ways to allocate bandwidth values to the LSPs, the Junos OS implementation supports only an equal-bandwidth allocation policy when normalization is done, wherein all the member LSPs are signaled or re-signaled with equal bandwidth.
- Number of LSPs—If an ingress router is configured to have a minimum number of LSPs, it maintains the minimum number of LSPs even if the demand can be satisfied with less than the minimum number of LSPs. In case the ingress router is unable to do constraint-based routing for computations on the sufficient number of LSPs or signal sufficient number of LSPs, the ingress router resorts to a number of fallback options.

By default, an incremental approach is supported as a fallback option (unless configured differently), where an ingress router makes attempts to bring up the sufficient number of LSPs, such that the new aggregate bandwidth exceeds the old aggregate bandwidth (and is as close to the desired demand as possible). The ingress router then load-balances traffic using the LSPs. The LSPs that could not be brought up are removed by the ingress router.

Supported Criteria

When a container LSP signals a member LSP, the member LSP gets signaled with minimum-signaling-bandwidth. Since each member LSP is configured with autobandwidth, between two normalization events, each LSP can undergo autobandwidth adjustment multiple times. As the traffic demand increases, the ingress router creates additional supplementary LSPs. All member LSPs are used for ECMP, so they should roughly have the same reserved bandwidth after normalization.

For example, if there are K LSPs signaled after normalization, each LSP is signaled with equal bandwidth B. The total aggregate bandwidth reserved is B.K, where B satisfies the following condition:

- Minimum signaling-bandwidth is less than or equal to B, which in turn is less than or equal to the maximum signaling-bandwidth

$$(\text{minimum-signaling-bandwidth} \leq B \leq \text{maximum-signaling-bandwidth})$$

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

$$(\text{minimum-bandwidth} \leq b_i \leq \text{maximum-bandwidth})$$

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**.

$$(\text{New-Aggr-Bw} > \text{Aggregate-maximum-bandwidth})$$

- **Relative trigger**—The **Current-Aggr-Bw** is compared with **New-Aggr-Bw** at the ingress routing device. LSP splitting is performed when the difference in the bandwidth amount is off by a threshold.

$$([1-a] \times \text{Current-Aggr-Bw} < \text{New-Aggr-Bw} < [1+a] \times \text{Current-Aggr-Bw}, \text{ where } 0 \leq a \leq 1)$$

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

IN THIS SECTION

- [Operational Overview | 740](#)
- [Operational Constraints | 740](#)
- [Merging Triggers | 741](#)

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] x Current-Aggr-Bw < New-Aggr-Bw < [1+a] x Current-Aggr-Bw, where 0 <= a <= 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

IN THIS SECTION

- [Operational Overview | 743](#)
- [Operational Constraints | 743](#)
- [Inter-Operation with Autobandwidth | 744](#)

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

IN THIS SECTION

- [Changes in Per-LSP Autobandwidth Adjustments | 744](#)
- [Changes in Traffic Growth | 746](#)
- [Computed Range and Configured Feasible Ranges | 746](#)

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

[Table 21 on page 744](#) illustrates how normalization splits and merges member LSPs as autobandwidth adjustments change per-LSP bandwidth with unconditional normalization.

Table 21: Normalization with Per-LSP Autobandwidth Adjustment Changes

Normalization Time	Current State	Events	Adjusted State
T0	No state.	Initialization	LSP-1 is signaled with bandwidth of 0.8G
T1	LSP-1 usage increases to 1.5G	<ul style="list-style-type: none"> • Multiple autobandwidth adjustments since T0 is possible. • The ingress router decides to split LSP-1 into two LSPs, and creates LSP-2. 	LSP-1 = 0.8G LSP-2 = 0.8G
T2	LSP-1 usage increase to 2G LSP-2 usage increases to 0.9G (within limits)	<ul style="list-style-type: none"> • Aggregate bandwidth is 2.9G, which exceeds aggregate splitting maximum of 2G. • The ingress router decides to split LSP-1 into three LSPs, and creates LSP-3. 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G

Table 21: 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"> Aggregate bandwidth is 3.5G with a maximum aggregate splitting of 3G. The ingress router decides to split LSP-1 into four LSPs, and creates LSP-4. 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G LSP-4 = 1G
T4	LSP-2 usage drops to 0.5G	<ul style="list-style-type: none"> Aggregate bandwidth is 3G. The ingress router decides to merge LSP-1 and removes LSP-4. 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G

Because autobandwidth is configured on a per-LSP basis, every time there is an autobandwidth adjustment, the ingress router re-signals each LSP with **Max Avg Bw**.

Another approach to handling the changes in per-LSP autobandwidth adjustments is to not allow individual LSPs to run autobandwidth on the ingress router, but to run autobandwidth in passive (monitor) mode. This way, sampling is done at every statistics interval for member LSPs only, and normalization is performed for the container LSP alone instead of acting on individual LSPs adjustment timer expiry.

As a result, the number of re-signaling attempts and bandwidth fluctuations for a given member LSP is reduced. Only the computed bandwidth-values per-member LSP is used by the ingress router to find an aggregate bandwidth to be used during normalization. Configuring autobandwidth adjustment followed by normalization (adjustments and normalization intervals are comparable) can lead to considerable overhead because of re-signaling.

Taking the same example, and applying the second approach, LSP-1 goes from 0.8G to 1.5G and then back to 0.8G. If the normalization timer is of the same order as the adjustment interval, the ingress router leaves LSP-1 alone with its original 0.8G and only signals LSP-2 with 0.8G. This helps achieve the final result of normalization, thus avoiding the extra signaling attempt on LSP-1 with 1.5G at adjustment timer expiry.

Because member LSPs always use equal bandwidth, any adjustment done on member LSPs is undone. The member LSPs are re-signaled with reduced bandwidth when compared to the reserved capacity in adjustment trigger with normalization trigger. Therefore, avoiding adjustment trigger for member LSPs might be useful assuming that normalization and adjustment intervals are of the same order.

NOTE: We recommend that the normalization timer be higher than the autobandwidth adjustment interval and regular optimization duration, as the traffic trends are observed at a longer time scale and normalization is performed one-to-three times per day. An LSP can undergo optimization for the following reasons:

- Normal optimization
- Autobandwidth adjustment
- Normalization

Changes in Traffic Growth

Table 22 on page 746 illustrates how normalization is performed when traffic grows in large factor.

Table 22: Normalization with Traffic Growth

Normalization Time	Current State	Events	Adjusted State
T0	No state		LSP-1 is signaled with bandwidth of 0.8G
T1	LSP-1 usage increase to 3G	<ul style="list-style-type: none"> • Aggregate usage exceeds maximum splitting bandwidth • The ingress router decides to split LSP-1, and creates two more supplementary LSPs 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G

Having fewer LSPs is preferred over signaling four LSPs each with 0.8G bandwidth, unless there is a constraint on the minimum number of LSPs.

Computed Range and Configured Feasible Ranges

When an ingress router is configured with the minimum and maximum number of LSPs, and per LSP splitting-bandwidth and merging-bandwidth values, the bandwidth thresholds are used for splitting and merging. For this, the number of LSPs (N) should satisfy the following constraints:

$$\text{minimum-member-lsps} \leq N \leq \text{maximum-member-lsps}$$

At the time of normalization, based on the aggregate demand X:

$$[X/\text{splitting-bandwidth}] \leq N \leq [X/\text{merging-bandwidth}]$$

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

IN THIS SECTION

- [NSR Support | 750](#)
- [IPG-FA Support | 751](#)
- [Static Route Support | 752](#)

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

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 23 on page 753 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 23: Applicability of RSVP LSPs Configuration to a Container LSP

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
adaptive (Default: non-adaptive)	Yes	Yes
admin-down	Yes	Yes
admin-group	Yes	Yes
admin-groups-except	Yes	Yes

Table 23: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
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 23: 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

Table 23: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
node-link-protection (All LSPs share same node-link protection mechanism)	Yes	Yes
oam	Yes	Yes
optimize-hold-dead-delay (All LSPs have same value)	Yes	Yes
optimize-switchover-delay (All LSPs have same value)	Yes	Yes
optimize-timer (All LSPs have same value)	Yes	Yes
p2mp	Yes	N/A
policing (Variable traffic)	Yes	No
preference	Yes	Yes
primary (All paths are primary)	Yes	No
random	Yes	Yes
record	Yes	Yes
retry-limit (Applicable to members)	Yes	Yes
retry-timer (Applicable to members)	Yes	Yes

Table 23: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
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
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 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
- Multichassis
- IPv6

RELATED DOCUMENTATION

[Example: Configuring Dynamic Bandwidth Management Using Container LSP | 759](#)

[Maximize Bandwidth Utilization with Juniper Networks TE++](#)

Example: Configuring Dynamic Bandwidth Management Using Container LSP

IN THIS SECTION

- [Requirements | 760](#)
- [Overview | 760](#)
- [Configuration | 761](#)
- [Verification | 772](#)

This example shows how to enable dynamic bandwidth management by configuring container label-switched paths (LSPs) that enable load balancing across multiple point-to-point member LSPs.

Requirements

This example uses the following hardware and software components:

- Five routers that can be a combination of M Series, MX Series, or T Series routers, out of which two routers are provider edge (PE) routers and three routers are provider (P) routers
- Junos OS Release 14.2 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
 - RSVP
 - MPLS
 - BGP
 - OSPF

Overview

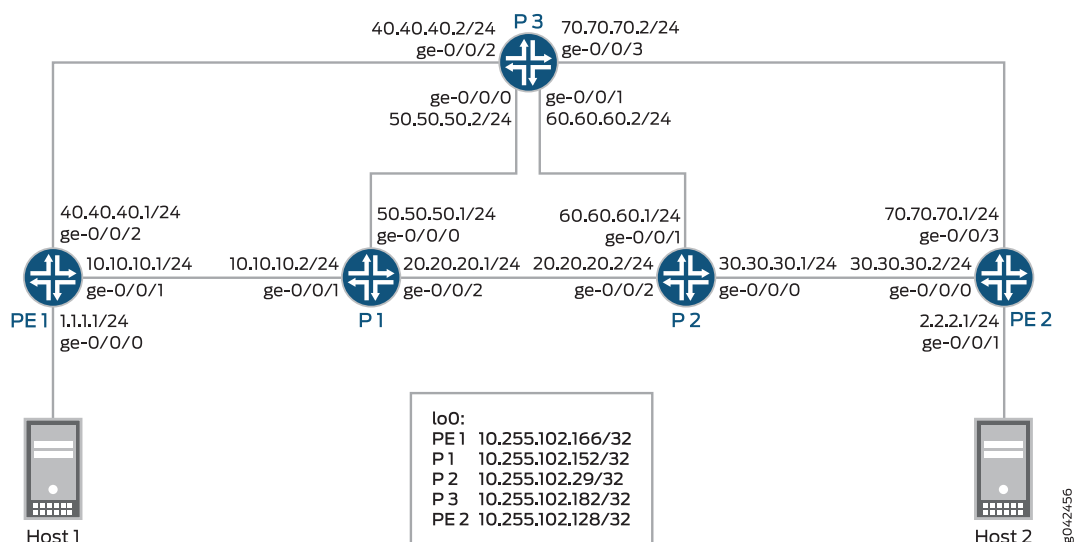
Starting with Junos OS Release 14.2, a new type of LSP, called a container LSP, is introduced to enable load balancing across multiple point-to-point LSPs. A container LSP includes one or more member LSPs between the same ingress and egress routing devices. The member LSPs are similar to independent point-to-point LSPs, and each member LSP takes a different path to the same destination and can be routed along a different IGP cost path.

A container LSP provides support for dynamic bandwidth management by enabling the ingress router to dynamically add and remove member LSPs through a process called LSP splitting and LSP merging, respectively, based on configuration and aggregate traffic. Besides addition and deletion, member LSPs can also be re-optimized with different bandwidth values in a make-before-break way.

Topology

[Figure 55 on page 761](#) is a sample topology configured with container LSPs.

Figure 55: Dynamic Bandwidth Management Using Container LSP



In this example, Routers PE1 and PE2 are the PE routers connected to hosts Host1 and Host2, respectively. The core routers, Routers P1, and P2, and P3 connect to the PE routers.

Configuration

CLI Quick Configuration

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

PE1

```

set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.1/24
set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 40.40.40.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.166/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.255.102.166
set routing-options autonomous-system 1234
set routing-options forwarding-table export pplib
set protocols rsvp preemption aggressive
  
```



```

set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls statistics file auto-bw
set protocols mpls statistics file size 10m
set protocols mpls statistics interval 10
set protocols mpls statistics auto-bandwidth
set protocols mpls label-switched-path PE1-to-PE2-template1 template
set protocols mpls label-switched-path PE1-to-PE2-template1 optimize-timer 30
set protocols mpls label-switched-path PE1-to-PE2-template1 link-protection
set protocols mpls label-switched-path PE1-to-PE2-template1 adaptive
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth adjust-interval 300
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth adjust-threshold 5
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth minimum-bandwidth
    10m
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth maximum-bandwidth
    10m
set protocols mpls container-label-switched-path PE1-PE2-container-100 label-switched-path-template
    PE1-to-PE2-template1
set protocols mpls container-label-switched-path PE1-PE2-container-100 to 10.255.102.128
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    maximum-member-lsps 20
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    minimum-member-lsps 2
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    splitting-bandwidth 40m
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    merging-bandwidth 6m
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    maximum-signaling-bandwidth 10m
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    minimum-signaling-bandwidth 10m
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    normalization normalize-interval 400
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    normalization failover-normalization
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    normalization normalization-retry-duration 20
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    normalization normalization-retry-limits 3

```



```

set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging sampling
  cut-off-threshold 1
set protocols mpls container-label-switched-path PE1-PE2-container-100 splitting-merging sampling
  use-percentile 90
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group to-PE2 type internal
set protocols bgp group to-PE2 local-address 10.255.102.166
set protocols bgp group to-PE2 family inet-vpn unicast
set protocols bgp group to-PE2 export direct
set protocols bgp group to-PE2 neighbor 10.255.102.128
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 metric 100
set policy-options policy-statement direct term 1 from protocol direct
set policy-options policy-statement direct term 1 then accept
set policy-options policy-statement pplb then load-balance per-packet
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-0/0/0.0
set routing-instances vpn1 route-distinguisher 10.255.102.166:1
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label

```

P1

```

set interfaces ge-0/0/0 unit 0 family inet address 50.50.50.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.152/32
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```



```
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 100
```

P2

```
set interfaces ge-0/0/0 unit 0 family inet address 30.30.30.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 60.60.60.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.29/32
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls statistics file auto_bw
set protocols mpls statistics file size 10m
set protocols mpls statistics interval 5
set protocols mpls statistics auto-bandwidth
set protocols mpls icmp-tunneling
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 metric 100
```

P3

```
set interfaces ge-0/0/0 unit 0 family inet address 50.50.50.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 60.60.60.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 40.40.40.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 70.70.70.2/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.182/32
set protocols rsvp interface all aggregate
```



```

set protocols rsvp interface fxp0.0 disable
set protocols mpls icmp-tunneling
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

PE2

```

set interfaces ge-0/0/0 unit 0 family inet address 30.30.30.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.1/24
set interfaces ge-0/0/3 unit 0 family inet address 70.70.70.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.128/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.255.102.128
set routing-options autonomous-system 1234
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group to-PE1 type internal
set protocols bgp group to-PE1 local-address 10.255.102.128
set protocols bgp group to-PE1 family inet-vpn unicast
set protocols bgp group to-PE1 neighbor 10.255.102.166
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-0/0/1.0
set routing-instances vpn1 route-distinguisher 10.255.102.128:1
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label

```

Step-by-Step Procedure

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

To configure Router PE1:

1. Configure the Router PE1 interfaces.

```
[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 1.1.1.1/24
user@PE1# set ge-0/0/1 unit 0 family inet address 10.10.10.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set ge-0/0/2 unit 0 family inet address 40.40.40.1/24
user@PE1# set ge-0/0/2 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 10.255.102.166/32
user@PE1# set lo0 unit 0 family mpls
```

2. Configure the router ID and autonomous system number for Router PE1.

```
[edit routing-options]
user@PE1# set router-id 10.255.102.166
user@PE1# set autonomous-system 1234
```

3. Enable the policy to load-balance traffic.

```
[edit routing-options]
user@PE1# set forwarding-table export pplb
```

4. Enable RSVP on all Router PE1 interfaces (excluding the management interface).

```
[edit protocols]
user@PE1# set rsvp preemption aggressive
user@PE1# set rsvp interface all aggregate
user@PE1# set rsvp interface fxp0.0 disable
user@PE1# set rsvp interface ge-0/0/1.0
user@PE1# set rsvp interface ge-0/0/2.0
```

5. Enable MPLS on all the interfaces of Router PE1 (excluding the management interface).

```
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
```


6. Configure the MPLS statistics parameters.

```
[edit protocols]
user@PE1# set mpls statistics file auto-bw
user@PE1# set mpls statistics file size 10m
user@PE1# set mpls statistics interval 10
user@PE1# set mpls statistics auto-bandwidth
```

7. Configure the label-switched path (LSP) template parameters.

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 template
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 optimize-timer 30
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 link-protection
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 adaptive
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth adjust-interval 300
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth adjust-threshold 5
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth minimum-bandwidth 10m
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth maximum-bandwidth 10m
```

8. Configure a container LSP between Router PE1 and Router PE2, and assign the PE1-to-PE2-template1 LSP template.

```
[edit protocols]
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 to 10.255.102.128
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 label-switched-path-template
    PE1-to-PE2-template1
```

9. Configure the container LSP parameters.

```
[edit protocols]
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    maximum-member-lsps 20
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    minimum-member-lsps 2
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    splitting-bandwidth 40m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    merging-bandwidth 6m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    maximum-signaling-bandwidth 10m
```



```

user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
minimum-signaling-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
normalize-interval 400
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
failover-normalization
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
normalization-retry-duration 20
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
normalization-retry-limits 3
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging sampling
cut-off-threshold 1
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging sampling
use-percentile 90

```

10. Configure the BGP group, and assign the local and neighbor IP addresses.

```

[edit protocols]
user@PE1# set bgp group to-PE2 type internal
user@PE1# set bgp group to-PE2 local-address 10.255.102.166
user@PE1# set bgp group to-PE2 neighbor 10.255.102.128
user@PE1# set bgp group to-PE2 family inet-vpn unicast
user@PE1# set bgp group to-PE2 export direct

```

11. Enable OSPF on all the interfaces of Router PE1 (excluding the management interface) along with traffic engineering capabilities.

```

[edit protocols]
user@PE1# set ospf traffic-engineering
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface ge-0/0/2.0 metric 100

```

12. Configure the policy statement to load-balance traffic.

```

[edit policy-options]
user@PE1# set policy-statement direct term 1 from protocol direct
user@PE1# set policy-statement direct term 1 then accept
user@PE1# set policy-statement pplb then load-balance per-packet

```

13. Configure a routing instance on Router PE1, and assign the routing instance interface.


```
[edit routing-instances]
user@PE1# set vpn1 instance-type vrf
user@PE1# set vpn1 interface ge-0/0/0.0
```

14. Configure the route distinguisher, vrf target, and vrf-table label values for the VRF routing instance.

```
[edit routing-instances]
user@PE1# set vpn1 route-distinguisher 10.255.102.166:1
user@PE1# set vpn1 vrf-target target:1:1
user@PE1# set vpn1 vrf-table-label
```

Results

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

```
user@PE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 1.1.1.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.10.10.1/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 40.40.40.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
```



```

    family inet {
        address 10.255.102.166/32;
    }
    family mpls;
}

```

user@PE1# show routing-options

```

router-id 10.255.102.166;
autonomous-system 1234;
forwarding-table {
    export pplb;
}

```

user@PE1# show protocols

```

rsvp {
    preemption aggressive;
    interface all {
        aggregate;
    }
    interface fxp0.0 {
        disable;
    }
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
}
mpls {
    statistics {
        file auto-bw size 10m;
        interval 10;
        auto-bandwidth;
    }
    label-switched-path PE1-to-PE2-template1 {
        template;
        optimize-timer 30;
        link-protection;
        adaptive;
        auto-bandwidth {
            adjust-interval 300;
            adjust-threshold 5;
            minimum-bandwidth 10m;
            maximum-bandwidth 10m;
        }
    }
}

```



```

}
container-label-switched-path PE1-PE2-container-100 {
    label-switched-path-template {
        PE1-to-PE2-template1;
    }
    to 10.255.102.128;
    splitting-merging {
        maximum-member-lsps 20;
        minimum-member-lsps 2;
        splitting-bandwidth 40m;
        merging-bandwidth 6m;
        maximum-signaling-bandwidth 10m;
        minimum-signaling-bandwidth 10m;
        normalization {
            normalize-interval 400;
            failover-normalization;
            normalization-retry-duration 20;
            normalization-retry-limits 3;
        }
        sampling {
            cut-off-threshold 1;
            use-percentile 90;
        }
    }
}
interface all;
interface fxp0.0 {
    disable;
}
}
bgp {
    group to-PE2 {
        type internal;
        local-address 10.255.102.166;
        family inet-vpn {
            unicast;
        }
        export direct;
        neighbor 10.255.102.128;
    }
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {

```



```

interface all;
interface fxp0.0 {
    disable;
}
interface ge-0/0/2.0 {
    metric 100;
}
}
}

```

user@PE1# show policy-options

```

policy-statement direct {
    term 1 {
        from protocol direct;
        then accept;
    }
}
policy-statement pplb {
    then load-balance {
        per-packet;
    }
}
}

```

user@PE1# show routing-instances

```

vpn1 {
    instance-type vrf;
    interface ge-0/0/0.0;
    route-distinguisher 10.255.102.166:1;
    vrf-target target:1:1;
    vrf-table-label;
}

```

Verification

IN THIS SECTION

- [Verifying the Container LSP Status Without Bandwidth | 773](#)
- [Verifying the Container LSP Status with Increased Bandwidth \(Before Normalization\) | 777](#)
- [Verifying the Container LSP Status with Increased Bandwidth \(After Normalization\) | 779](#)
- [Verifying the Container LSP Splitting Process | 783](#)

- [Verifying the Container LSP Statistics | 784](#)
- [Verifying the Container LSP Status with Decreased Bandwidth \(Before Normalization\) | 785](#)
- [Verifying the Container LSP Status with Decreased Bandwidth \(After Normalization\) | 785](#)
- [Verifying the Container LSP Merging Process | 786](#)
- [Verifying Failover Normalization | 787](#)
- [Verifying Incremental Normalization | 788](#)

Confirm that the configuration is working properly.

Verifying the Container LSP Status Without Bandwidth

Purpose

Verify the status of the container LSP.

Action

From operational mode, run the **show mpls container-lsp extensive** command.

```
user@PE1> show mpls container-lsp extensive
```

```
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 2
Normalization
  Min LSPs: 2, Max LSPs: 20
  Aggregate bandwidth: 20Mbps, Sampled Aggregate bandwidth: 0bps
  NormalizeTimer: 400 secs, NormalizeThreshold: 10%
  Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
  BW: 6Mbps
  Mode: incremental-normalization, failover-normalization
  Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
  Normalization in 143 second(s)
    36 Jun  5 04:12:17.497 Clear history and statistics: on container
    (PE1-PE2-container-100)
    35 Jun  5 04:12:17.497 Avoid normalization: not needed with bandwidth 0 bps
    34 Jun  5 04:05:37.484 Clear history and statistics: on container
    (PE1-PE2-container-100)
    33 Jun  5 04:05:37.484 Avoid normalization: not needed with bandwidth 0 bps
    32 Jun  5 03:58:57.470 Clear history and statistics: on container
    (PE1-PE2-container-100)
    31 Jun  5 03:58:57.470 Avoid normalization: not needed with bandwidth 0 bps
```



```

30 Jun  5 03:52:17.455 Clear history and statistics: on container
(PE1-PE2-container-100)
29 Jun  5 03:52:17.455 Avoid normalization: not needed with bandwidth 0 bps
28 Jun  5 03:45:37.440 Clear history and statistics: on container
(PE1-PE2-container-100)
27 Jun  5 03:45:37.440 Avoid normalization: not needed with bandwidth 0 bps
26 Jun  5 03:38:59.013 Normalization complete: container (PE1-PE2-container-100)
with 2 members
25 Jun  5 03:38:57.423 Delete member LSPs: PE1-PE2-container-100-3 through
PE1-PE2-container-100-7
24 Jun  5 03:38:57.423 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 7
23 Jun  5 03:38:57.423 Normalize: normalization with aggregate bandwidth 0 bps

22 Jun  5 03:32:19.019 Normalization complete: container (PE1-PE2-container-100)
with 7 members
21 Jun  5 03:32:17.404 Clear history and statistics: on container
(PE1-PE2-container-100)
20 Jun  5 03:32:17.403 Normalize: container (PE1-PE2-container-100) into 7
members - each with bandwidth 10000000 bps
19 Jun  5 03:32:17.403 Normalize: normalization with aggregate bandwidth 62914560
bps
18 Jun  5 03:32:17.403 Normalize: normalizaton with 62914560 bps
17 Jun  5 03:32:09.219 Normalization complete: container (PE1-PE2-container-100)
with 7 members
16 Jun  5 03:32:07.600 Clear history and statistics: on container
(PE1-PE2-container-100)
15 Jun  5 03:32:07.600 Normalize: container (PE1-PE2-container-100) into 7
members - each with bandwidth 10000000 bps
14 Jun  5 03:32:07.599 Normalize: normalization with aggregate bandwidth 62914560
bps
13 Jun  5 03:32:07.599 Normalize: normalizaton with 62914560 bps
12 Jun  5 03:26:57.295 Clear history and statistics: on container
(PE1-PE2-container-100)
11 Jun  5 03:26:57.295 Avoid normalization: not needed with bandwidth 0 bps
10 Jun  5 03:20:18.297 Normalization complete: container (PE1-PE2-container-100)
with 2 members
9 Jun  5 03:20:17.281 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
8 Jun  5 03:20:17.281 Normalize: normalization with aggregate bandwidth 0 bps

7 Jun  5 03:17:43.218 Clear history and statistics: on container
(PE1-PE2-container-100)
6 Jun  5 03:17:43.218 Avoid normalization: not needed with bandwidth 0 bps

```



```

    5 Jun  5 03:17:43.212 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-2
    4 Jun  5 03:17:43.212 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-1
    3 Jun  5 03:12:47.323 Normalization complete: container (PE1-PE2-container-100)
with 2 members
    2 Jun  5 03:12:16.555 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
    1 Jun  5 03:12:16.555 Normalize: normalization with aggregate bandwidth 0 bps

10.255.102.128
  From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-1

  ActivePath: (primary)
  LSPTtype: Dynamic Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 10Mbps, MaxBW: 10Mbps
  AdjustTimer: 300 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 12 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.10.10.2 20.20.20.2 30.30.30.2
  17 Jun  5 03:38:59.013 Make-before-break: Switched to new instance
  16 Jun  5 03:38:58.003 Record Route:  10.10.10.2 20.20.20.2 30.30.30.2
  15 Jun  5 03:38:58.003 Up
  14 Jun  5 03:38:57.423 Originate make-before-break call
  13 Jun  5 03:38:57.423 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
  12 Jun  5 03:33:30.400 CSPF: computation result ignored, new path no benefit
  11 Jun  5 03:32:17.403 Pending old path instance deletion
  10 Jun  5 03:32:09.218 Make-before-break: Switched to new instance
   9 Jun  5 03:32:08.202 Record Route:  10.10.10.2 20.20.20.2 30.30.30.2
   8 Jun  5 03:32:08.202 Up
   7 Jun  5 03:32:07.603 Originate make-before-break call

```



```

    6 Jun  5 03:32:07.603 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
    5 Jun  5 03:20:18.278 Selected as active path
    4 Jun  5 03:20:18.277 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
    3 Jun  5 03:20:18.277 Up
    2 Jun  5 03:20:17.281 Originate Call
    1 Jun  5 03:20:17.281 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
Created: Thu Jun  5 03:20:16 2014

10.255.102.128
From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-2

ActivePath: (primary)
LSPTtype: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs
Max AvgBW util: 0bps, Bandwidth Adjustment in 76 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.10.10.2 20.20.20.2 30.30.30.2
    17 Jun  5 03:38:59.013 Make-before-break: Switched to new instance
    16 Jun  5 03:38:58.002 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
    15 Jun  5 03:38:58.002 Up
    14 Jun  5 03:38:57.423 Originate make-before-break call
    13 Jun  5 03:38:57.423 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
    12 Jun  5 03:33:26.189 CSPF: computation result ignored, new path no benefit
    11 Jun  5 03:32:17.403 Pending old path instance deletion
    10 Jun  5 03:32:09.219 Make-before-break: Switched to new instance
    9 Jun  5 03:32:08.204 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
    8 Jun  5 03:32:08.204 Up
    7 Jun  5 03:32:07.603 Originate make-before-break call

```



```

    6 Jun  5 03:32:07.603 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
    5 Jun  5 03:20:18.297 Selected as active path
    4 Jun  5 03:20:18.295 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
    3 Jun  5 03:20:18.295 Up
    2 Jun  5 03:20:17.281 Originate Call
    1 Jun  5 03:20:17.281 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
Created: Thu Jun  5 03:20:16 2014
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The container LSP is established between Routers PE1 and PE2.

Verifying the Container LSP Status with Increased Bandwidth (Before Normalization)

Purpose

Verify the status of the container LSP with increased bandwidth before normalization happens.

Action

From operational mode, run the **show mpls container-lsp extensive** command.

```
user@PE1> show mpls container-lsp extensive
```

```

Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 2
Normalization
  Min LSPs: 2, Max LSPs: 20
  Aggregate bandwidth: 20Mbps, Sampled Aggregate bandwidth: 42.6984Mbps
  NormalizeTimer: 400 secs, NormalizeThreshold: 10%
  Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
BW: 6Mbps
  Mode: incremental-normalization, failover-normalization
  Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
  Normalization in 321 second(s)
    3 Jun  5 21:22:34.731 Normalization complete: container (PE1-PE2-container-100)
with 2 members

```



```

    2 Jun  5 21:22:15.503 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
    1 Jun  5 21:22:15.503 Normalize: normalization with aggregate bandwidth 0 bps

10.255.102.128
  From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-1

  ActivePath: (primary)
  Link protection desired
  LSPTtype: Dynamic Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 10Mbps, MaxBW: 10Mbps
  AdjustTimer: 300 secs AdjustThreshold: 5%
  Max AvgBW util: 23.9893Mbps, Bandwidth Adjustment in 221 second(s).
  Overflow limit: 0, Overflow sample count: 6
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps
  OptimizeTimer: 30
  SmartOptimizeTimer: 180
  Reoptimization in 9 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.255.102.166(flag=0x20) 10.10.10.2(Label=303440)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302144) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
  From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-2

  ActivePath: (primary)
  Link protection desired
  LSPTtype: Dynamic Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 10Mbps, MaxBW: 10Mbps
  AdjustTimer: 300 secs AdjustThreshold: 5%
  Max AvgBW util: 22.1438Mbps, Bandwidth Adjustment in 221 second(s).
  Overflow limit: 0, Overflow sample count: 6

```



```

Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps
  OptimizeTimer: 30
  SmartOptimizeTimer: 180
  Reoptimization in 9 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
      10.255.102.166(flag=0x20) 10.10.10.2(Label=303456)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302160) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

Total 2 displayed, Up 2, Down 0

```

Meaning

Because normalization has not happened, the member LSP count remains at 2.

Verifying the Container LSP Status with Increased Bandwidth (After Normalization)

Purpose

Verify the status of the container LSP with increased bandwidth after normalization happens.

Action

From operational mode, run the **show mpls container-lsp extensive** command.

```
user@PE1> show mpls container-lsp extensive
```

```

Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 5
Normalization
  Min LSPs: 2, Max LSPs: 20
  Aggregate bandwidth: 50Mbps, Sampled Aggregate bandwidth: 45.8873Mbps
  NormalizeTimer: 400 secs, NormalizeThreshold: 10%
  Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
  BW: 6Mbps
  Mode: incremental-normalization, failover-normalization
  Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
  Normalization in 169 second(s)
    7 Jun  5 21:29:02.921 Normalization complete: container (PE1-PE2-container-100)

```



```

with 5 members
  6 Jun  5 21:28:55.505 Clear history and statistics: on container
(PE1-PE2-container-100)
  5 Jun  5 21:28:55.505 Normalize: container (PE1-PE2-container-100) into 5
members - each with bandwidth 10000000 bps
  4 Jun  5 21:28:55.504 Normalize: normalization with aggregate bandwidth 45281580
bps
  3 Jun  5 21:22:34.731 Normalization complete: container (PE1-PE2-container-100)
with 2 members
  2 Jun  5 21:22:15.503 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
  1 Jun  5 21:22:15.503 Normalize: normalization with aggregate bandwidth 0 bps

10.255.102.128
  From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-1

  ActivePath: (primary)
  Link protection desired
  LSPTtype: Dynamic Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 10Mbps, MaxBW: 10Mbps
  AdjustTimer: 300 secs AdjustThreshold: 5%
  Max AvgBW util: 11.0724Mbps, Bandwidth Adjustment in 129 second(s).
  Overflow limit: 0, Overflow sample count: 1
  Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps
  OptimizeTimer: 30
  SmartOptimizeTimer: 180
  Reoptimization in 12 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.255.102.166(flag=0x20) 10.10.10.2(Label=303488)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302224) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
  From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-2

```



```

ActivePath: (primary)
Link protection desired
LSPTtype: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs AdjustThreshold: 5%
Max AvgBW util: 8.50751Mbps, Bandwidth Adjustment in 189 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 11, Underflow Max AvgBW: 8.50751Mbps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps
  OptimizeTimer: 30
  SmartOptimizeTimer: 180
  Reoptimization in 6 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.255.102.166(flag=0x20) 10.10.10.2(Label=303504)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302240) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
  From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-3

ActivePath: (primary)
Link protection desired
LSPTtype: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs AdjustThreshold: 5%
Max AvgBW util: 9.59422Mbps, Bandwidth Adjustment in 249 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 5, Underflow Max AvgBW: 9.59422Mbps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps

```



```

    OptimizeTimer: 30
    SmartOptimizeTimer: 180
    Reoptimization in 25 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
        10.255.102.166(flag=0x20) 10.10.10.2(Label=303472)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302176) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
    From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-4

    ActivePath: (primary)
    Link protection desired
    LSPTtype: Dynamic Configured, Penultimate hop popping
    LoadBalance: Random
    Autobandwidth
    MinBW: 10Mbps, MaxBW: 10Mbps
    AdjustTimer: 300 secs AdjustThreshold: 5%
    Max AvgBW util: 9.16169Mbps, Bandwidth Adjustment in 9 second(s).
    Overflow limit: 0, Overflow sample count: 0
    Underflow limit: 0, Underflow sample count: 29, Underflow Max AvgBW: 9.16169Mbps

    Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                                State: Up
    Priorities: 7 0
    Bandwidth: 10Mbps
    OptimizeTimer: 30
    SmartOptimizeTimer: 180
    Reoptimization in 1 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
        10.255.102.166(flag=0x20) 10.10.10.2(Label=303520)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302192) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
    From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-5

    ActivePath: (primary)

```



```

Link protection desired
LSPTtype: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs AdjustThreshold: 5%
Max AvgBW util: 8.39908Mbps, Bandwidth Adjustment in 69 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 23, Underflow Max AvgBW: 8.39908Mbps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  Bandwidth: 10Mbps
  OptimizeTimer: 30
  SmartOptimizeTimer: 180
  Reoptimization in 17 second(s).
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
      10.255.102.166(flag=0x20) 10.10.10.2(Label=303536)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302208) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)
Total 5 displayed, Up 5, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

At the expiry of the normalization timer, the container LSP is split into five member LSPs, each with 10 Mbps (minimum and maximum signaling bandwidth). As a result, the aggregate bandwidth is 50 Mbps.

Verifying the Container LSP Splitting Process

Purpose

Verify the container LSP splitting process after normalization happens.

Action

From operational mode, run the **show route 2.2.2** command.


```
user@PE1> show route 2.2.2
```

```

vpn1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.0/24          *[BGP/170] 00:12:14, localpref 100, from 10.255.102.128
                    AS path: I, validation-state: unverified
>to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-1
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-2
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-3
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-4
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-5

```

Meaning

After LSP splitting, Router PE1 has injected the forwarding adjacency.

Verifying the Container LSP Statistics

Purpose

Verify the container LSP statistics after normalization happens.

Action

From operational mode, run the **show mpls container-lsp statistics** command.

```
user@PE1> show mpls container-lsp statistics
```

```

Ingress LSP: 1 sessions
Container LSP name                                State      Member LSP count
PE1-PE2-container-100                            Up          5
To          From          State      Packets      Bytes LSPname
10.255.102.128  10.255.102.166 Up    15166271    2062612856
PE1-PE2-container-100-1
10.255.102.128  10.255.102.166 Up    12289912    1671428032
PE1-PE2-container-100-2
10.255.102.128  10.255.102.166 Up    13866911    1885899896
PE1-PE2-container-100-3
10.255.102.128  10.255.102.166 Up    12558707    1707984152
PE1-PE2-container-100-4
10.255.102.128  10.255.102.166 Up    11512151    1565652536
PE1-PE2-container-100-5

```

Meaning

Traffic is load-balanced across the newly created member LSPs.

Verifying the Container LSP Status with Decreased Bandwidth (Before Normalization)

Purpose

Verify the status of the container LSP with decreased bandwidth before normalization happens.

Action

From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
```

```
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 5
Normalization
  Min LSPs: 2, Max LSPs: 20
  Aggregate bandwidth: 50Mbps, Sampled Aggregate bandwidth: 2.0215Mbps
  NormalizeTimer: 400 secs, NormalizeThreshold: 10%
  Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
  BW: 6Mbps
  Mode: incremental-normalization, failover-normalization
  Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
  Normalization in 384 second(s)
---Output truncated---
```

Meaning

Because normalization has not happened, the member LSP count remains at 5.

Verifying the Container LSP Status with Decreased Bandwidth (After Normalization)

Purpose

Verify the status of the container LSP with decreased bandwidth after normalization happens.

Action

From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
```

```
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 2
Normalization
  Min LSPs: 2, Max LSPs: 20
  Aggregate bandwidth: 20Mbps, Sampled Aggregate bandwidth: 0bps
```



```

NormalizeTimer: 400 secs, NormalizeThreshold: 10%
Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
BW: 6Mbps
Mode: incremental-normalization, failover-normalization
Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
Normalization in 397 second(s)
 22 Jun  5 22:30:37.094 Clear history and statistics: on container
(PE1-PE2-container-100)
 21 Jun  5 22:30:37.094 Delete member LSPs: PE1-PE2-container-100-3 through
PE1-PE2-container-100-5
 20 Jun  5 22:30:37.090 Normalize: container (PE1-PE2-container-100) into 2
members - each with bandwidth 10000000 bps
 19 Jun  5 22:30:37.090 Normalize: normalization with aggregate bandwidth 2037595
bps
 18 Jun  5 22:30:37.090 Normalize: normalizaton with 2037595 bps
---Output truncated---

```

Meaning

At the expiry of the normalization timer, the container LSP merging takes place because there is an overall reduction in bandwidth. The member LSPs are merged, and the member LSP count is 2 after normalization.

Verifying the Container LSP Merging Process

Purpose

Verify the container LSP splitting process after normalization happens.

Action

From operational mode, run the **show route 2.2.2** command.

```
user@PE1> show route 2.2.2
```

```

vpn1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.0/24          *[BGP/170] 01:09:45, localpref 100, from 10.255.102.128
                    AS path: I, validation-state: unverified
                    > to 10.10.10.2 via ge-0/0/1.0, label-switched-path
PE1-PE2-container-100-1
                    to 10.10.10.2 via ge-0/0/1.0, label-switched-path
PE1-PE2-container-100-2

```

Meaning

After LSP merging, Router PE1 has deleted the merged member LSPs.

Verifying Failover Normalization

Purpose

Verify load redistribution when traffic is sent at 35 Mbps and the link between Routers P1 and P2 is disabled. Arrival of PathErr on link failure triggers immediate normalization.

To enable failover normalization, include the **failover-normalization** configuration statement at the [edit protocols mpls container-label-switched-path container-lsp-name splitting-merging normalization] hierarchy level.

Action

From operational mode, run the **show mpls container-lsp** command.

```
user@PE1> show mpls container-lsp
```

```
Ingress LSP: 1 sessions
Container LSP name          State      Member LSP count
PE1-PE2-container-100      Up          2
To          From          State Rt P    ActivePath    LSPname
10.255.102.128  10.255.102.166  Up      0  *
PE1-PE2-container-100-1
10.255.102.128  10.255.102.166  Up      0  *
PE1-PE2-container-100-2
Total 2 displayed, Up 2, Down 0
```

After the ge-0/0/2 link between Routers P1 and P2 goes down, normalization is immediately triggered.

From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
```

```
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 4
Normalization
  Min LSPs: 2, Max LSPs: 20
  Aggregate bandwidth: 40Mbps, Sampled Aggregate bandwidth: 34.5538Mbps
  NormalizeTimer: 3000 secs, NormalizeThreshold: 10%
  Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
  BW: 6Mbps
  Mode: incremental-normalization, failover-normalization
  Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
  Normalization in 2970 second(s)
```



```

11 Jun 5 19:28:27.564 Normalization complete: container (PE1-PE2-container-100)
with 4 members
10 Jun 5 19:28:20.315 Normalize: container (PE1-PE2-container-100) received PathErr on member
PE1-PE2-container-100-2[2 times]
9 Jun 5 19:28:20.315 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-1[2 times]
8 Jun 5 19:28:20.311 Clear history and statistics: on container
(PE1-PE2-container-100)
7 Jun 5 19:28:20.311 Normalize: container (PE1-PE2-container-100) into 4 members
- each with bandwidth 10000000 bps
6 Jun 5 19:28:20.311 Normalize: normalization with aggregate bandwidth 33665020
bps
5 Jun 5 19:28:20.308 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-2
4 Jun 5 19:28:20.308 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-1
3 Jun 5 19:27:48.574 Normalization complete: container (PE1-PE2-container-100)
with 2 members
2 Jun 5 19:27:28.644 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
1 Jun 5 19:27:28.644 Normalize: normalization with aggregate bandwidth 0 bps
----Output truncated----

```

Meaning

Arrival of PathErr message on link failure triggers immediate normalization.

Verifying Incremental Normalization

Purpose

Verify incremental normalization when enough bandwidth is not available.

On Router PE1, the RSVP interfaces static bandwidth is restricted to 22 Mbps each.

Action

From operational mode, run the **show rsvp interface** command.

```
user@PE1> show rsvp interface
```

```

RSVP interface: 4 active

```

Interface	State	Active resv	Subscr- iption	Static BW	Available BW	Reserved BW	Highwater mark
ge-0/0/2.0	Up	0	100%	22Mbps	22Mbps	0bps	21.4031Mbps
ge-0/0/1.0	Up	2	100%	22Mbps	12Mbps	10Mbps	21.7011Mbps

Before normalization happens:

From operational mode, run the **show mpls container-lsp** command.

```
user@PE1> show mpls container-lsp
```

```
Ingress LSP: 1 sessions
Container LSP name          State      Member LSP count
PE1-PE2-container-100      Up          2
To          From          State Rt P    ActivePath  LSPname
10.255.102.128  10.255.102.166  Up      0  *
PE1-PE2-container-100-1
10.255.102.128  10.255.102.166  Up      0  *
PE1-PE2-container-100-2
```

After normalization happens:

From operational mode, run the **show mpls container-lsp** command.

```
user@PE1> show mpls container-lsp
```

```
Ingress LSP: 1 sessions
Container LSP name          State      Member LSP count
PE1-PE2-container-100      Up          7
To          From          State Rt P    ActivePath  LSPname
10.255.102.128  10.255.102.166  Up      0  *    PE1-PE2-container-100-1
10.255.102.128  10.255.102.166  Up      0  *    PE1-PE2-container-100-2
10.255.102.128  10.255.102.166  Up      0  *    PE1-PE2-container-100-3
10.255.102.128  10.255.102.166  Up      0  *    PE1-PE2-container-100-4
10.255.102.128  10.255.102.166  Up      0  *    PE1-PE2-container-100-5
10.255.102.128  10.255.102.166  Up      0  *    PE1-PE2-container-100-6
10.255.102.128  0.0.0.0          Dn      0  -    PE1-PE2-container-100-7
Total 7 displayed, Up 6, Down 1
```

From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
```

```
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 7
Normalization
Min LSPs: 2, Max LSPs: 10
Aggregate bandwidth: 40.8326Mbps, Sampled Aggregate bandwidth: 50.129Mbps
```



```

NormalizeTimer: 9000 secs, NormalizeThreshold: 10%
Max Signaling BW: 10Mbps, Min Signaling BW: 5Mbps, Splitting BW: 40Mbps, Merging
BW: 5Mbps
Mode: incremental-normalization, failover-normalization
Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
Normalization in 8072 second(s)
10 Jun 5 18:40:17.812 Normalization complete: container (PE1-PE2-container-100)
with 7 members, retry-limit reached
9 Jun 5 18:40:08.028 Normalize: container (PE1-PE2-container-100) for target
member count 7, member bandwidth 6805439 bps
8 Jun 5 18:39:58.301 Normalize: container (PE1-PE2-container-100) for target
member count 6, member bandwidth 7939679 bps
7 Jun 5 18:39:48.470 Clear history and statistics: on container
(PE1-PE2-container-100)
6 Jun 5 18:39:48.470 Normalize: container (PE1-PE2-container-100) into 5 members
- each with bandwidth 9527615 bps
5 Jun 5 18:39:48.469 Normalize: normalization with aggregate bandwidth 47638076
bps
4 Jun 5 18:39:48.469 Normalize: normalizaton with 47638076 bps
3 Jun 5 18:39:09.471 Normalization complete: container (PE1-PE2-container-100)
with 2 members
2 Jun 5 18:38:59.822 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 5000000bps, member count 0
1 Jun 5 18:38:59.822 Normalize: normalization with aggregate bandwidth 0 bps

```

Meaning

After normalization, the aggregate bandwidth after three retries is 40.8326 Mbps.

RELATED DOCUMENTATION

[Dynamic Bandwidth Management Using Container LSP Overview](#) | 729

Configuring Dynamic Bandwidth Management Using Container LSP

You can configure a container LSP to enable load balancing across multiple point-to-point LSPs dynamically. A container LSP includes one or more member LSPs between the same ingress and egress routing devices. The member LSPs are similar to independent point-to-point LSPs, and each member LSP takes a different path to the same destination and can be routed along a different IGP cost path.

A container LSP provides support for dynamic bandwidth management by enabling the ingress router to dynamically add and remove member LSPs through a process called LSP splitting and LSP merging, respectively, based on configuration and aggregate traffic. Besides addition and deletion, member LSPs can also be re-optimized with different bandwidth values in a make-before-break way.

Before you begin:

1. Configure the device interfaces.
2. Configure the device router ID and autonomous system number.
3. Configure the following protocols:
 - RSVP
 - BGP

Configure a BGP group to peer device with remote provider edge (PE) device.
 - OSPF

Enable traffic engineering capabilities.
4. Configure a VRF routing instance.

To configure the PE device:

1. Enable MPLS on all the interfaces (excluding the management interface).

```
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
```

2. Configure the MPLS statistics parameters.

```
[edit protocols]
user@PE1# set mpls statistics file file-name
user@PE1# set mpls statistics file size size
user@PE1# set mpls statistics interval seconds
user@PE1# set mpls statistics auto-bandwidth
```


3. Configure the label-switched path (LSP) template parameters.

```
[edit protocols]
user@PE1# set mpls label-switched-path template-name template
user@PE1# set mpls label-switched-path template-name optimize-timer seconds
user@PE1# set mpls label-switched-path template-name link-protection
user@PE1# set mpls label-switched-path template-name adaptive
user@PE1# set mpls label-switched-path template-name auto-bandwidth adjust-interval seconds
user@PE1# set mpls label-switched-path template-name auto-bandwidth adjust-threshold seconds
user@PE1# set mpls label-switched-path template-name auto-bandwidth minimum-bandwidth mbps
user@PE1# set mpls label-switched-path template-name auto-bandwidth maximum-bandwidth mbps
```

4. Configure a container LSP between the two PE routers, and assign the LSP template.

```
[edit protocols]
user@PE1# set mpls container-label-switched-path container-lsp-name to remote-PE-ip-address
user@PE1# set mpls container-label-switched-path container-lsp-name label-switched-path-template
template-name
```

5. Configure the container LSP parameters.

```
[edit protocols]
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging
    maximum-member-lsps number
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging
    minimum-member-lsps number
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging splitting-bandwidth
    mbps
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging merging-bandwidth
    mbps
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging
    maximum-signaling-bandwidth mbps
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging
    minimum-signaling-bandwidth mbps
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging normalization
    normalize-interval seconds
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging normalization
    failover-normalization
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging normalization
    normalization-retry-duration seconds
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging normalization
    normalization-retry-limits number
```



```

user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging sampling
cut-off-threshold number
user@PE1# set mpls container-label-switched-path container-lsp-name splitting-merging sampling
use-percentile number

```

6. Configure the policy statement to load-balance traffic.

```

[edit policy-options]
user@PE1# set policy-statement first-policy-name term 1 from protocol direct
user@PE1# set policy-statement first-policy-name term 1 then accept
user@PE1# set policy-statement second-policy-name then load-balance per-packet

```

NOTE: The policy to load-balance traffic should be assigned to the forwarding table configuration under the [edit routing-options] hierarchy level.

```

user@PE1# set forwarding-table export pplb

```

7. Verify and commit the configuration.

For example:

```

[edit protocols]
user@PE1# set rsvp preemption aggressive
user@PE1# set rsvp interface all aggregate
user@PE1# set rsvp interface fxp0.0 disable
user@PE1# set rsvp interface ge-0/0/1.0
user@PE1# set rsvp interface ge-0/0/2.0
user@PE1# set mpls statistics file auto-bw
user@PE1# set mpls statistics file size 10m
user@PE1# set mpls statistics interval 10
user@PE1# set mpls statistics auto-bandwidth
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 template
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 optimize-timer 30
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 link-protection
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 adaptive
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth adjust-interval 300
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth adjust-threshold 5
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth minimum-bandwidth 10m
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth maximum-bandwidth 10m

```



```

user@PE1# set mpls label-switched-path PE1-PE2-template-1 template
user@PE1# set mpls label-switched-path PE1-PE2-template-1 auto-bandwidth adjust-interval 8000
user@PE1# set mpls label-switched-path PE1-PE2-template-1 auto-bandwidth minimum-bandwidth 5m
user@PE1# set mpls label-switched-path PE1-PE2-template-1 auto-bandwidth maximum-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 label-switched-path-template
    PE1-to-PE2-template1
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 to 10.255.102.128
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    maximum-member-lsps 20
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    minimum-member-lsps 2
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    splitting-bandwidth 40m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    merging-bandwidth 6m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    maximum-signaling-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging
    minimum-signaling-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
    normalize-interval 400
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
    failover-normalization
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
    normalization-retry-duration 20
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging normalization
    normalization-retry-limits 3
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging sampling
    cut-off-threshold 1
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 splitting-merging sampling
    use-percentile 90
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set bgp group to-PE2 type internal
user@PE1# set bgp group to-PE2 local-address 10.255.102.166
user@PE1# set bgp group to-PE2 family inet-vpn unicast
user@PE1# set bgp group to-PE2 export direct
user@PE1# set bgp group to-PE2 neighbor 10.255.102.128
user@PE1# set ospf traffic-engineering
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface ge-0/0/2.0 metric 100

```

[edit policy-options]


```
user@PE1# set policy-statement direct term 1 from protocol direct
user@PE1# set policy-statement direct term 1 then accept
user@PE1# set policy-statement pplb then load-balance per-packet
```

```
[edit]
user@PE1# commit
commit complete
```

RELATED DOCUMENTATION

[Dynamic Bandwidth Management Using Container LSP Overview | 729](#)

[Example: Configuring Dynamic Bandwidth Management Using Container LSP | 759](#)

Configuring Pop-and-Forward LSPs

IN THIS CHAPTER

- [Pop-and-Forward LSP Configuration | 797](#)

Pop-and-Forward LSP Configuration

IN THIS SECTION

- [Benefits of RSVP-TE Pop-and-Forward LSP Tunnels | 798](#)
- [Pop-and-Forward LSP Tunnel Terminology | 798](#)
- [Pop-and-Forward LSP Tunnel Label and Signaling | 799](#)
- [Pop-and-Forward LSP Tunnel Label Stacking | 800](#)
- [Pop-and-Forward LSP Tunnel Link Protection | 801](#)
- [RSVP-TE Pop-and-Forward LSP Tunnel Supported and Unsupported Features | 802](#)

Pop-and-forward LSPs introduces the notion of pre-installed per traffic engineering link pop labels that are shared by RSVP-TE LSPs that traverse these links and significantly reducing the required forwarding plane state . A transit label-switching router (LSR) allocates a unique pop label per traffic engineering link with a forwarding action to pop the label and forward the packet over that traffic engineering link should the label appear at the top of the packet. These pop labels are sent back in the RESV message of the LSP at each LSR and further recorded in the record route object (RRO). The label stack is constructed from the recorded labels in the RRO and pushed by the ingress label edge router (LER), as each transit hop performs a pop-and-forward action on its label. The pop-and-forward tunnels enhances the RSVP-TE control plane feature benefits with the simplicity of the shared MPLS forwarding plane.

Benefits of RSVP-TE Pop-and-Forward LSP Tunnels

- **Scaling advantage of RSVP-TE**—Any platform-specific label space limit on an LSR is prevented from being a constraint to the control plane scaling on that interface.
- **Reduced forwarding plane state**—The transit labels on a traffic engineering link are shared across RSVP-TE tunnels traversing the link, and are used independent of the ingress and egress devices of the LSPs, thereby significantly reducing the required forwarding plane state.
- **Reduced transit data plane state**—Because the pop labels are allocated per traffic engineering link and shared across LSPs, the total label state in the forwarding plane is reduced to a function of the number of RSVP neighbors on that interface.
- **Faster LSP setup time**—The forwarding plane state is not programmed during the LSP setup and teardown. As a result, the control plane need not wait sequentially at each hop for the forwarding plane to be programmed prior to sending the label upstream in the RESV message, resulting in reduced LSP setup time.
- **Backward compatibility**—This allows backward compatibility with transit LSRs that provide regular labels in RESV messages. Labels can be mixed across transit hops in a single MPLS RSVP-TE LSP. Certain LSRs can use traffic engineering link labels and others can use regular labels. The ingress can construct a label stack appropriately based on what type of label is recorded from every transit LSR.

Pop-and-Forward LSP Tunnel Terminology

The following terminology is used in the implementation of RSVP-TE pop-and-forward LSP tunnels:

- **Pop label**—An incoming label at an LSR that is popped and forwarded over a specific traffic-engineering link to a neighbor.
- **Swap label**—An incoming label at an LSR that is swapped to an outgoing label and forwarded over a specific downstream traffic engineering link.
- **Delegation label**—An incoming label at an LSR that is popped. A new set of labels is pushed before the packet is forwarded.
- **Delegation hop**— A transit hop that allocates a delegation label.
- **Application label depth (AppLD)**—Maximum number of application or service labels (for example, VPN, LDP, or IPv6 explicit-null labels) that can be beneath the RSVP transport labels. It is configured on a per-node basis, and is equally applicable for all LSPs, and is neither signaled nor advertised.
- **Outbound label depth (OutLD)**—Maximum number of labels that can be pushed before a packet is forwarded. This is local to the node, and is neither signaled nor advertised.
- **Additional transport label depth (AddTLD)**—Maximum number of other transport labels that can be added (for example, bypass label). This is a per-LSP parameter that is neither signaled or advertised. The

value is discerned by checking if the LSP has been signaled with link protection (AddTLD=1) or without link protection (AddTLD=0).

- **Effective transport label depth (ETLD)**— Number of transport labels that the LSP hop can potentially send to its downstream hop. This value is signaled per LSP in the hop attributes subobject. The hop attributes subobject is added to the record route object (RRO) in the path message.

Pop-and-Forward LSP Tunnel Label and Signaling


Every traffic engineering link is allocated a pop label that is installed in the mpls.0 routing table with a forwarding action to pop the label and forward the packet over the traffic engineering link to the downstream neighbor of the RSVP-TE tunnel.

For pop-and-forward LSP tunnels, the pop label for the traffic engineering link is allocated when the first RESV message for a pop-and-forward transit LSP arrives over that traffic engineering link. This is done to avoid preallocating pop labels and installing them in networks where pop-and-forward LSPs are not configured.

NOTE: For the pop-and-forward LSP tunnels to function effectively, we recommend that you configure the **maximum-labels** statement on all the interfaces in the RSVP-TE network.

Figure 56 on page 799 displays pop labels at all interfaces for neighboring devices.

Figure 56: Pop-and-Forward LSP Tunnel Labels



There are two pop-and-forward LSP tunnels—T1 and T2. Tunnel T1 is from Device A to Device E on path A-B-C-D-E. Tunnel T2 is from Device F to Device E on path F-B-C-D-E. Both the tunnels, T1 and T2, share the same traffic engineering links B-C, C-D, and D-E.

As RSVP-TE signals the setup of the pop-and-forward tunnel T1, the LSR D receives the RESV message from the egress E. Device D checks the next-hop traffic engineering link (D-E) and provides the pop label (250) in the RESV message for the tunnel. The label is sent in the label object and is also recorded in the label subobject (with the pop label bit set) carried in the RRO. Similarly, Device C provides the pop label

(200) for the next-hop traffic engineering link C-D and Device B provides the pop label (150) for the next-hop traffic engineering link B-C. For the tunnel T2, the transit LSRs provide the same pop labels as described for tunnel T1.

Both the label edge routers (LERs), Device A and Device F, push the same stack of labels [150(top), 200, 250] for tunnels T1 and T2, respectively. The recorded labels in the RRO are used by the ingress LER to construct a stack of labels.

The pop-and-forward LSP tunnel labels are compatible with transit interfaces that use swap labels. Labels can be mixed across transit hops in a single MPLS RSVP-TE LSP, where certain LSRs can use pop labels and others can use swap labels. The ingress device constructs the appropriate label stack based on the label type recorded from every transit LSR.

Pop-and-Forward LSP Tunnel Label Stacking

Construction of Label Stack at the Ingress

The ingress LER checks the type of label received from each transit hop as recorded in the RRO in the RESV message and generates the appropriate label stack to use for the pop-and-forward tunnel.

The following logic is used by the ingress LER while constructing the label stack:

- Each RRO label subobject is processed starting with the label subobject from the first downstream hop.
- Any label provided by the first downstream hop is always pushed on the label stack. If the label type is a pop label, then any label from the succeeding downstream hop is also pushed on the constructed label stack.
- If the label type is a swap label, then any label from the succeeding downstream hop is not pushed on the constructed label stack.

Auto-Delegation of Label Stack

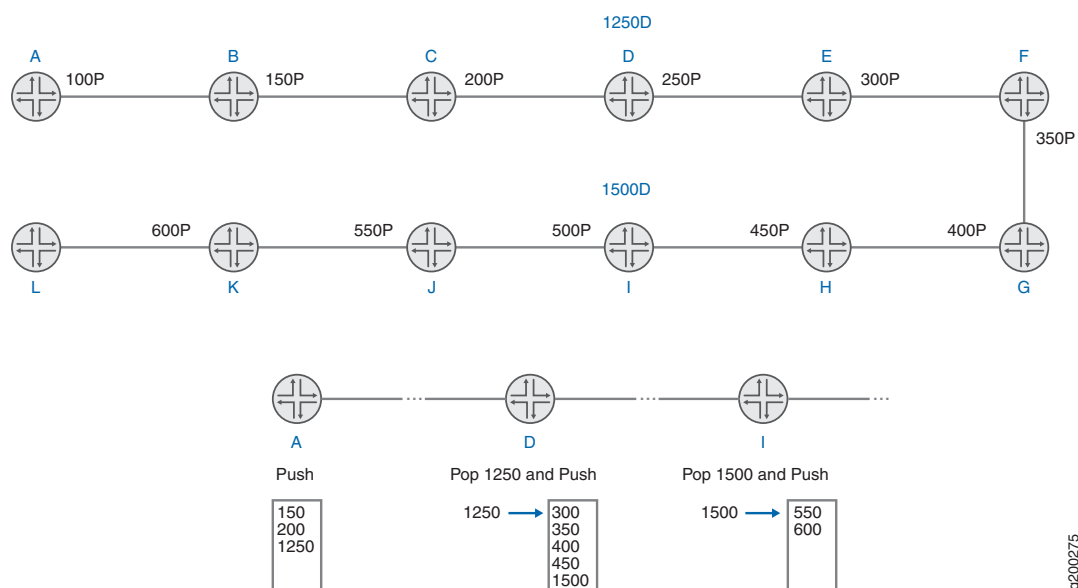
The ingress device runs the Constrained Shortest Path First (CSPF) to compute the path, and if the hop length is greater than the OutLD-AppLD-AddTLD, the ingress device cannot impose the entire label stack to reach the egress device.

When requesting RSVP-TE to signal the path, the ingress device always requests autodelegation for the LSP, where one or more transit hops automatically select themselves as delegation hops to push the label stack to reach the next delegation hop. Junos OS uses an algorithm based on the received Effective Transport Label-Stack Depth (ETLD), that each transit executes to decide whether it should autoselect itself as a delegation hop. This algorithm is based on the section on ETLD in the Internet draft draft-ietf-mpls-rsvp-shared-labels-00.txt (expires September 11 2017), *Signaling RSVP-TE Tunnels on a Shared MPLS Forwarding Plane*.

The label stack imposed by the ingress device delivers the packet up to the first delegation hop. Each delegation hop's label stack also includes the delegation label of the next delegation hop at the bottom of the stack.

Figure 57 on page 801 displays labels at every device interface, where Device D and Device I are delegation hops, [Label] P is the pop label, and [Label] D is the delegation label. The RSVP-TE pop-and-forward LSP tunnel is A-B-C-D-E-F-G-H-I-J-K-L. Delegation label 1250 represents (300, 350, 400, 450, 1500); Delegation label 1500 represents (550, 600).

Figure 57: Pop-and-Forward LSP Tunnel Pop and Delegation Labels



In this approach, for the tunnel, the ingress LER Device A pushes (150, 200, 1250). At LSR Device D, the delegation label 1250 gets popped and labels 300, 350, 400, 450, and 1500 get pushed. At LSR Device I, the delegation label 1500 gets popped and the remaining set of labels (550, 600) get pushed. In Junos OS, the pop and push action occurs as a swap to the bottom label of the outgoing stack and push the remaining labels.


A delegation label and the LSP segment that it covers can be shared by multiple pop-and-forward LSPs. A LSP delegation segment consist of an ordered set of hops (IP addresses and labels) as seen in the RESV RRO. The delegation label (and the segment that it covers) is not owned by a particular LSP, but can be shared. When all LSPs using a delegation label are deleted, the delegation label (and route) is deleted.

Pop-and-Forward LSP Tunnel Link Protection

To provide link protection at a point of local repair (PLR) with a pop-and-forward data plane, the LSR allocates a separate pop label for the traffic engineering link that is used for the RSVP-TE tunnels that request link protection from the ingress device. No signaling extensions are required to support link protection for the RSVP-TE tunnels over the pop-and-forward data plane.

Figure 58 on page 802 displays pop labels at every device interface; labels marked with P are pop labels that offer link protection for the traffic-engineering link.

Figure 58: Pop-and Forward LSP Tunnel Link Protection



At each LSR, link-protected pop labels can be allocated for each traffic engineering link, and a link-protecting facility bypass LSP (which is not a pop-and-forward LSP, but rather a normal bypass LSP) can be created to protect the traffic engineering link. These labels can be sent in the RESV message by the LSR for LSPs requesting link protection over the specific traffic engineering link. Because the facility bypass terminates at the next hop (merge point), the incoming pop label on the packet at the PLR is what the merge point expects.

For example, LSR Device B can install a facility bypass LSP for the link-protected pop label 151. When the traffic engineering link B-C is up, LSR Device B pops 151 and sends the packet to C. If the traffic engineering link B-C is down, the LSR can pop 151 and send the packet through the facility backup to Device C.

RSVP-TE Pop-and-Forward LSP Tunnel Supported and Unsupported Features

Junos OS supports the following features with RSVP-TE pop-and-forward LSP tunnels:

- Pop labels per RSVP neighbor for unprotected LSP.
- Pop labels per RSVP neighbor for LSPs requesting link protection using facility bypass
- Autodelegation of LSP segment.
- Mixed label mode, where certain transit LSRs do not support pop-and-forward LSP tunnels
- LSP ping and traceroute
- All existing CSPF constraint.
- Load balancing of traffic between pop-and-forward LSPs and regular point-to-point RSVP-TE LSP.
- Autobandwidth, LDP tunneling, and TE++ container LSP.
- Aggregated Ethernet interface.
- Virtual platforms support, such as Juniper Networks vMX Virtual Router.
- 64-bit support
- Logical systems

Junos OS does not support the following functionality for RSVP-TE pop-and-forward LSP tunnels:

- Node link protection
- Detour protection for MPLS fast reroute
- Point-to-multipoint LSPs.
- Switch-away LSP.
- Generalized MPLS (GMPLS) LSPs (including bidirectional LSPs, associated LSPs, VLAN user-to-network interface [UNI] and so on)
- IP Flow Information Export (protocol) (IPFIX) inline flow sampling for MPLS template
- RFC 3813, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)*
- IPv4 Explicit-null (Inserting label 0 at the bottom of the label stack is not supported. If there are service labels beneath the RSVP-TE pop-and-forward label stack, because the penultimate hop for the LSP copies the EXP value to the service label, this can allow continuity of class of service (CoS) across the MPLS forwarding plane).
- Ultimate-hop popping (UHP)
- Graceful Routing Engine switchover (GRES)
- Nonstop active routing (NSR)

4

PART

MPLS Signalling Protocols

Configuring RSVP | **807**

Configuring LDP | **917**

Configuring RSVP

IN THIS CHAPTER

- [RSVP Overview | 809](#)
- [RSVP Operation Overview | 809](#)
- [Understanding the RSVP Signaling Protocol | 810](#)
- [RSVP-TE protocol extensions for FRR | 813](#)
- [Junos OS RSVP Protocol Implementation | 816](#)
- [RSVP Authentication | 816](#)
- [Supported RSVP Standards | 816](#)
- [RSVP and IGP Hello Packets and Timers | 818](#)
- [RSVP Message Types | 819](#)
- [Understanding RSVP Automatic Mesh | 819](#)
- [Path Messages | 820](#)
- [Resv Messages | 821](#)
- [PathTear Messages | 821](#)
- [ResvTear Messages | 821](#)
- [PathErr Messages | 822](#)
- [ResvErr Messages | 822](#)
- [ResvConfirm Messages | 822](#)
- [RSVP Reservation Styles | 822](#)
- [RSVP Refresh Reduction | 823](#)
- [MTU Signaling in RSVP | 824](#)
- [How the Correct MTU Is Signaled in RSVP | 825](#)
- [Determining an Outgoing MTU Value | 826](#)
- [MTU Signaling in RSVP Limitations | 826](#)
- [Minimum RSVP Configuration | 827](#)
- [Configuring RSVP and MPLS | 827](#)
- [Configuring RSVP Interfaces | 829](#)
- [Configuring RSVP Node-ID Hellos | 835](#)
- [Example: Configuring RSVP-Signaled LSPs | 836](#)

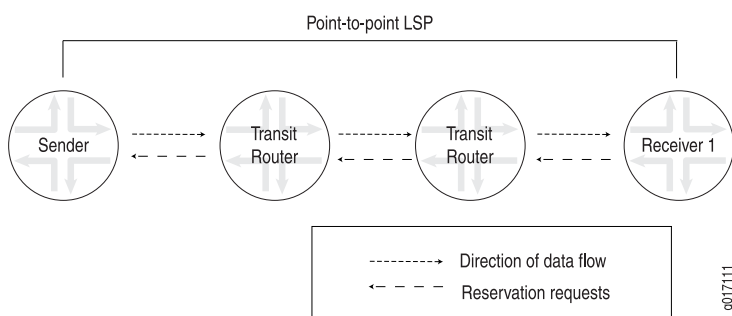
- [Example: Configuring RSVP Automatic Mesh | 841](#)
- [Configuring Hello Acknowledgments for Nonsession RSVP Neighbors | 846](#)
- [Switching LSPs Away from a Network Node | 846](#)
- [Configuring RSVP Setup Protection | 848](#)
- [Configuring Load Balancing Across RSVP LSPs | 848](#)
- [Configuring RSVP Automatic Mesh | 850](#)
- [Configuring Timers for RSVP Refresh Messages | 851](#)
- [Preempting RSVP Sessions | 852](#)
- [Configuring MTU Signaling in RSVP | 853](#)
- [Configuring Ultimate-Hop Popping for LSPs | 854](#)
- [Configuring RSVP to Pop the Label on the Ultimate-Hop Router | 858](#)
- [Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs | 859](#)
- [Tracing RSVP Protocol Traffic | 860](#)
- [Link Protection | 864](#)
- [Multiple Bypass LSPs for Link Protection | 865](#)
- [Node Protection | 865](#)
- [Fast Reroute, Node Protection, and Link Protection | 867](#)
- [Configuring Link Protection on Interfaces Used by LSPs | 871](#)
- [Configuring Node Protection or Link Protection for LSPs | 879](#)
- [Configuring Inter-AS Node and Link Protection | 880](#)
- [RSVP Graceful Restart | 881](#)
- [RSVP Graceful Restart Standard | 881](#)
- [RSVP Graceful Restart Terminology | 881](#)
- [RSVP Graceful Restart Operation | 882](#)
- [Processing the Restart Cap Object | 883](#)
- [Configuring RSVP Graceful Restart | 883](#)
- [RSVP LSP Tunnels Overview | 885](#)
- [Example: RSVP LSP Tunnel Configuration | 888](#)
- [Configuring Link Management Protocol Peers | 912](#)
- [Configuring Link Management Protocol Traffic Engineering Links | 913](#)
- [Configuring Peer Interfaces in OSPF and RSVP | 914](#)
- [Defining Label-Switched Paths for the FA-LSP | 915](#)
- [Establishing FA-LSP Path Information | 915](#)
- [Option: Tearing Down RSVP LSPs Gracefully | 916](#)

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 59 on page 809](#).
- 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 packets that can be sent over RSVP-signaled LSPs.

Figure 59: RSVP Reservation Request and Data Flow



RSVP Operation Overview

RSVP creates independent sessions to handle each data flow. A session is identified by a combination of the destination address, an optional destination port, and a protocol. Within a session, there can be one or more senders. Each sender is identified by a combination of its source address and source port. An out-of-band mechanism, such as a session announcement protocol or human communication, is used to communicate the session identifier to all senders and receivers.

A typical RSVP session involves the following sequence of events:

1. A potential sender starts sending RSVP path messages to the session address.
2. A receiver, wanting to join the session, registers itself if necessary. For example, a receiver in a multicast application would register itself with IGMP.
3. The receiver receives the path messages.
4. The receiver sends appropriate Resv messages toward the sender. These messages carry a flow descriptor, which is used by routers along the path to make reservations in their link-layer media.
5. The sender receives the Resv message and then starts sending application data.

This sequence of events is not necessarily strictly synchronized. For example, receivers can register themselves before receiving path messages from the sender, and application data can flow before the sender receives Resv messages. Application data that is delivered before the actual reservation contained in the Resv message typically is treated as best-effort, non-real-time traffic with no CoS guarantee.

Understanding the RSVP Signaling Protocol

IN THIS SECTION

- [RSVP Fundamentals | 811](#)
- [Bandwidth Reservation Requirement | 811](#)
- [Explicit Route Objects | 811](#)
- [Constrained Shortest Path First | 812](#)
- [Link Coloring | 813](#)

RSVP is a signaling protocol that handles bandwidth allocation and true traffic engineering across an MPLS network. Like LDP, RSVP uses discovery messages and advertisements to exchange LSP path information between all hosts. However, RSVP also includes a set of features that control the flow of traffic through an MPLS network. Whereas LDP is restricted to using the configured IGP's shortest path as the transit path through the network, RSVP uses a combination of the Constrained Shortest Path First (CSPF) algorithm and Explicit Route Objects (EROs) to determine how traffic is routed through the network.

Basic RSVP sessions are established in exactly the same way that LDP sessions are established. By configuring both MPLS and RSVP on the appropriate transit interfaces, you enable the exchange of RSVP packets and the establishment of LSPs. However, RSVP also lets you configure link authentication, explicit LSP paths, and link coloring.

This topic contains the following sections:

RSVP Fundamentals

RSVP uses unidirectional and simplex flows through the network to perform its function. The inbound router initiates an RSVP path message and sends it downstream to the outbound router. The path message contains information about the resources needed for the connection. Each router along the path begins to maintain information about a reservation.

When the path message reaches the outbound router, resource reservation begins. The outbound router sends a reservation message upstream to the inbound router. Each router along the path receives the reservation message and sends it upstream, following the path of the original path message. When the inbound router receives the reservation message, the unidirectional network path is established.

The established path remains open as long as the RSVP session is active. The session is maintained by the transmission of additional path and reservation messages that report the session state every 30 seconds. If a router does not receive the maintenance messages for 3 minutes, it terminates the RSVP session and reroutes the LSP through another active router.

Bandwidth Reservation Requirement

When a bandwidth reservation is configured, reservation messages propagate the bandwidth value throughout the LSP. Routers must reserve the bandwidth specified across the link for the particular LSP. If the total bandwidth reservation exceeds the available bandwidth for a particular LSP segment, the LSP is rerouted through another LSR. If no segments can support the bandwidth reservation, LSP setup fails and the RSVP session is not established.

Explicit Route Objects

Explicit Route Objects (EROs) limit LSP routing to a specified list of LSRs. By default, RSVP messages follow a path that is determined by the network IGP's shortest path. However, in the presence of a configured ERO, the RSVP messages follow the path specified.


EROs consist of two types of instructions: loose hops and strict hops. When a loose hop is configured, it identifies one or more transit LSRs through which the LSP must be routed. The network IGP determines the exact route from the inbound router to the first loose hop, or from one loose hop to the next. The loose hop specifies only that a particular LSR be included in the LSP.

When a strict hop is configured, it identifies an exact path through which the LSP must be routed. Strict-hop EROs specify the exact order of the routers through which the RSVP messages are sent.

You can configure loose-hop and strict-hop EROs simultaneously. In this case, the IGP determines the route between loose hops, and the strict-hop configuration specifies the exact path for particular LSP path segments.

Figure 60 on page 812 shows a typical RSVP-signaled LSP that uses EROs.

Figure 60: Typical RSVP-Signaled LSP with EROs



In the topology shown in Figure 60 on page 812, traffic is routed from Host C1 to Host C2. The LSP can pass through Routers R4 or Router R7. To force the LSP to use R4, you can set up either a loose-hop or strict-hop ERO that specifies R4 as a hop in the LSP. To force a specific path through Router R4, R3, and R6, configure a strict-hop ERO through the three LSRs.

Constrained Shortest Path First

Whereas IGP uses the Shortest Path First (SPF) algorithm to determine how traffic is routed within a network, RSVP uses the Constrained Shortest Path First (CSPF) algorithm to calculate traffic paths that are subject to the following constraints:

- LSP attributes—Administrative groups such as link coloring, bandwidth requirements, and EROs
- Link attributes—Colors on a particular link and available bandwidth

These constraints are maintained in the traffic engineering database (TED). The database provides CSPF with up-to-date topology information, the current reservable bandwidth of links, and the link colors.

In determining which path to select, CSPF follows these rules:

- Computes LSPs one at a time, beginning with the highest priority LSP—the one with the lowest setup priority value. Among LSPs of equal priority, CSPF starts with those that have the highest bandwidth requirement.
- Prunes the traffic engineering database of links that are not full duplex and do not have sufficient reservable bandwidth.

- If the LSP configuration includes the **include** statement, prunes all links that do not share any included colors.
- If the LSP configuration includes the **exclude** statement, prunes all links that contain excluded colors. If a link does not have a color, it is accepted.
- Finds the shortest path toward the LSP's outbound router, taking into account any EROs. For example, if the path must pass through Router A, two separate SPF algorithms are computed: one from the inbound router to Router A and one from Router A to the outbound router.
- If several paths have equal cost, chooses the one with a last-hop address the same as the LSP's destination.
- If several equal-cost paths remain, selects the path with the least number of hops.
- If several equal-cost paths remain, applies CSPF load-balancing rules configured on the LSP.

Link Coloring

RSVP allows you to configure administrative groups for CSPF path selection. An administrative group is typically named with a color, assigned a numeric value, and applied to the RSVP interface for the appropriate link. Lower numbers indicate higher priority.

After configuring the administrative group, you can either exclude, include, or ignore links of that color in the TED:

- If you exclude a particular color, all segments with an administrative group of that color are excluded from CSPF path selection.
- If you include a particular color, only those segments with the appropriate color are selected.
- If you neither exclude nor include the color, the metrics associated with the administrative groups and applied on the particular segments determine the path cost for that segment.

The LSP with the lowest total path cost is selected into the TED.

RELATED DOCUMENTATION

[MPLS Traffic Engineering and Signaling Protocols Overview](#) | 1134

[Example: Configuring RSVP-Signaled LSPs](#) | 836

RSVP-TE protocol extensions for FRR

Starting with Junos OS Release 16.1, RSVP Traffic Engineering (TE) protocol extensions to support Refresh-interval Independent RSVP (RI-RSVP) defined RFC 8370 for fast reroute (FRR) facility protection

were introduced to allow greater scalability of label-switched paths (LSPs) faster convergence times and decrease RSVP signaling message overhead from periodic refreshes. Junos RSVP-TE runs in enhanced FRR aka RI-RSVP mode by default that includes protocol extensions to support RI-RSVP for FRR facility bypass originally specified in RFC 4090.

The RI-RSVP protocol extensions implemented in Junos are fully backward compatible. In mixed environments, where a subset of LSPs traverse nodes that do not include this feature, Junos RSVP-TE running in enhanced FRR mode will automatically turn off the new protocol extensions in its signaling exchanges with nodes that do not support the new extensions.

As part of enhanced FRR profile, a number of changes were made and new defaults adopted. These are listed here.

- RSVP-TE runs “enhanced” FRR, or RI-RSVP mode, by default, which includes enhancements to facilitate scaled up scenarios. These new protocol extensions can be disabled on a router with the **no-enhanced-frr-bypass** command.
- RSVP refresh reduction extensions defined in RFC 2961 are enabled by default. The user can disable them with the **no-reliable** command (see below).

NOTE: Node-id based Hellos are enabled by default as enhanced FRR or RI-RSVP extensions require the exchange of Node-id based Hello messages. Node-id based Hellos can be disabled with **node-hello** command. As refresh-reduction extensions and node-id based Hello messages are essential for enhanced FRR or RI-RSVP extensions, disabling either of them will automatically turn off enhanced FRR extensions on the node.

- The default refresh time for RSVP messages has increased from 30 seconds to 20 minutes.
- The default value for **keep-multiplier**, which is 3, is applied to the new default refresh time.
- Local reversion is disabled by default. The existing CLI configuration for node Hellos, **[edit protocols rsvp node-hello]**, is still available but the action is redundant. If enabled, a message is displayed to indicate that the configuration is not necessary in conjunction with enhanced FRR.
- The existing commands to disable message reliability are now used to disable RSVP refresh reduction. To revert back to the previous default enabling refresh reduction, use the **delete** version of the following commands:
 - Set **no-reliable** on a given interface to disable FRR scalability enhancements automatically for all LSPs traversing the interface. Likewise, to run RSVP-TE without FRR facility protection enhancements, and without refresh-reduction, disable refresh reduction on each interface. Use one of the commands shown here:
 - **[edit protocols rsvp interface name no-reliable]**

- [edit logical-systems *name* protocols rsvp interface *name* no-reliable]
- Graceful restart and nonstop active routing (NSR) are not supported while the LSP undergoes local repair or while the LSP is refreshed during backup LSP signaling. This limitation exists already in the implementation because GR or NSR switchover during FRR local-repair makes for multiple failure scenario.
- The following operational commands have been updated to include new information about the new procedures introduced for the RSVP TE protocol extensions for FRR facility protection.
 - **show rsvp version** displays whether enhanced FRR procedures are enabled.
 - **show rsvp neighbor detail** displays whether enhanced FRR procedures are enabled on the neighbor.
 - **show rsvp interface detail** displays conditional PathTear statistics.
 - **show rsvp statistics** displays sent and received statistics for conditional PathTear, along with other statistics.
 - **show rsvp session extensive** displays whether the node is a merge point, and if it is, shows the Point of Local Repair (PLR) address.
- The previous CLI configuration options for enabling or disabling message bundling have been deprecated (the existing configurations are accepted, but a warning is displayed in the show configuration output). These commands are the following:
 - [edit protocols rsvp interface *name* aggregate]
 - [edit logical-systems *name* protocols rsvp interface *name* aggregate]
 - [edit protocols rsvp interface *name* no-aggregate]
 - [edit logical-systems *name* protocols rsvp interface *name* no-aggregate]
- The following CLI configuration options have been made redundant by the current changes (the existing configurations are accepted, but a warning is displayed in the show configuration output):
 - [edit protocols rsvp interface *name* reliable]
 - [edit logical-systems *name* protocols rsvp interface *name* reliable]

Release History Table

Release	Description
16.1	Starting with Junos OS Release 16.1, RSVP Traffic Engineering (TE) protocol extensions to support Refresh-interval Independent RSVP (RI-RSVP) defined RFC 8370 for fast reroute (FRR) facility protection were introduced to allow greater scalability of label-switched paths (LSPs) faster convergence times and decrease RSVP signaling message overhead from periodic refreshes.

RELATED DOCUMENTATION

[Configuring RSVP Refresh Reduction | 829](#)

[Configuring RSVP Node-ID Hellos | 835](#)

Junos OS RSVP Protocol Implementation

The Junos implementation of RSVP supports RSVP version 1. The software includes support for all mandatory objects and RSVP message types, and supports message integrity and node authentications through the Integrity object.

The primary purpose of the Junos RSVP software is to support dynamic signaling within MPLS label-switched paths (LSPs). Supporting resource reservations over the Internet is only a secondary purpose of the Junos OS implementation. Since supporting resource reservations is secondary, the Junos RSVP software does not support the following features:

- IP multicasting sessions.
- Traffic control. The software cannot make resource reservations for real-time video or audio sessions.

With regard to the protocol mechanism, packet processing, and RSVP objects supported, the Junos OS implementation of the software is interoperable with other RSVP implementations.

RSVP Authentication

The Junos OS supports both the RSVP authentication style described in RFC 2747 (allowing for multivendor compatibility) and the RSVP authentication style described in Internet draft draft-ietf-rsvp-md5-03.txt. The Junos OS uses the authentication style described in Internet draft draft-ietf-rsvp-md5-08.txt by default. If the router receives an RFC 2747-style RSVP authentication from a neighbor, it switches to this style of authentication for that neighbor. The RSVP authentication style for each neighboring router is determined separately.

Supported RSVP Standards

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

- RFC 2205, *Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification*
- RFC 2210, *The Use of RSVP with IETF Integrated Services*

- RFC 2211, *Specification of the Controlled-Load Network Element Service*
- RFC 2212, *Specification of Guaranteed Quality of Service*
- RFC 2215, *General Characterization Parameters for Integrated Service Network Elements*
- RFC 2745, *RSVP Diagnostic Messages*
- RFC 2747, *RSVP Cryptographic Authentication* (updated by RFC 3097)
- RFC 2750, *RSVP Extensions for Policy Control* (RFC is not supported. Fully compliant with devices that support this RFC).
- RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
- RFC 3097, *RSVP Cryptographic Authentication—Updated Message Type Value*
- RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

The Null Service Object for maximum transmission unit (MTU) signaling in RSVP is not supported.

- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

Only Section 9, “Fault Handling,” is supported.

- RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*
- RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
(OSPF extensions can carry traffic engineering information over unnumbered links.)
- RFC 4558, *Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement*
- RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

The RRO node ID subobject is for use in inter-AS link and node protection configurations.

- RFC 4875, *Extensions to RSVP-TE for Point-to-Multipoint TE LSPs*
- RFC 5420, *Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)*

Only the LSP_ATTRIBUTES object is supported.

- RFC 7570, *Label Switched Path (LSP) Attribute in the Explicit Route Object (ERO)*
- RFC 8370, *Techniques to Improve the Scalability of RSVP-TE Deployments*
- draft-ietf-mpls-ri-rsvp-frr-05, *Refresh Interval Independent FRR Facility Protection*
- draft-ietf-mpls-rsvp-shared-labels-09, *Signaling RSVP-TE tunnels on a shared MPLS forwarding plane*

The following RFCs do not define standards, but provide information about RSVP and related technologies. The IETF classifies them variously as “Experimental” or “Informational.”

- RFC 2209, *Resource ReSerVation Protocol (RSVP)—Version 1 Message Processing Rules*
- RFC 2216, *Network Element Service Specification Template*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RELATED DOCUMENTATION

[Supported GMPLS Standards | 1356](#)

[Supported LDP Standards | 923](#)

[Supported MPLS Standards | 20](#)

[Supported PCEP Standards](#)

[Accessing Standards Documents on the Internet](#)

RSVP and IGP Hello Packets and Timers

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello packets are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

In the Junos OS, RSVP typically relies on IGP hello packet detection to check for node failures. Configuring a short time for the IS-IS or OSPF hello timers allows these protocols to detect node failures quickly. When the node fails or a node failure is detected, a path error message is generated, and although the RSVP session goes down, the IGP neighbors remain up.

RSVP hellos can be relied on when the IGP does not recognize a particular neighbor (for example, if IGP is not enabled on the interface) or if the IGP is RIP (not IS-IS or OSPF). Also, the equipment of other vendors might be configured to monitor RSVP sessions based on RSVP hello packets. This equipment might also take an RSVP session down due to a loss of RSVP hello packets.

We do not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

OSPF and IS-IS have infrastructure to manage rapid hello message sending and receiving reliably, even if the routing protocols or some other process are straining the processing capability of the router. Under

the same circumstances, RSVP hellos might time out prematurely even though the neighbor is functioning normally.

RSVP Message Types

RSVP uses the following types of messages to establish and remove paths for data flows, establish and remove reservation information, confirm the establishment of reservations, and report errors:

- [Path Messages on page 820](#)
- [Resv Messages on page 821](#)
- [PathTear Messages on page 821](#)
- [ResvTear Messages on page 821](#)
- [PathErr Messages on page 822](#)
- [ResvErr Messages on page 822](#)
- [ResvConfirm Messages on page 822](#)

Understanding RSVP Automatic Mesh

When adding sites to BGP and MPLS VPNs that use RSVP signaling, more configuration is needed to add provider edge (PE) routers than is needed to add customer edge (CE) devices. RSVP automatic mesh helps to reduce this configuration burden.

Service providers often use BGP and MPLS VPNs to efficiently scale the network while delivering revenue-generating services. In these environments, BGP is used to distribute the VPN routing information across the service provider's network, while MPLS is used to forward that VPN traffic from one VPN site to another. BGP and MPLS VPNs are based on a peer model. To add a new CE device (site) to an existing VPN, you need to configure the CE router at the new site and the PE router connected to the CE router. You do not have to modify the configuration of all of the other PE routers participating in the VPN. The other PE routers automatically learn about the routes associated with the new site, a process called automatic discovery (AD).

The requirements are a bit different if you need to add a new PE router to the network. A BGP and MPLS VPN requires that the BGP session be fully meshed and that there also be a full mesh of PE router-to-PE router MPLS label-switched paths (LSPs) between all of the PE routers in the network. When you add a new PE router to the network, all of the existing PE routers must be reconfigured to peer with the new PE router. Much of the configuration effort can be reduced if you configure BGP route reflectors (mitigating the full mesh requirement for BGP) and if you configure (LDP) as the signaling protocol for MPLS.

However, if you need to add a new PE router to a network configured with a full mesh of RSVP-signaled LSPs, you must reconfigure each of the PE routers to have a peer relationship with the new PE router. You can configure RSVP automatic mesh to address this particular operational scenario. When you enable RSVP automatic mesh, RSVP LSPs are dynamically created between a new PE router and the existing PE routers, eliminating the need to reconfigure all of the PE routers manually. For dynamic LSP creation to function properly, BGP must be configured to exchange routes between all of the participating PE routers. If two BGP peers do not exchange routes, it is not possible to configure a dynamic LSP between them. The local router's inet.0 routing table must include a labeled route to each potential IBGP next-hop (future potential PE routers or LSP destinations).

RSVP includes numerous capabilities that are not available in LDP, including fast reroute, end-point control, and link management. RSVP automatic mesh helps to reduce the operation and maintenance requirements for RSVP, making it possible to deploy RSVP in larger and more complicated networks.

Every PE router can reach every other PE router in the network because this information is distributed by the IGP. A PE router can set up a point-to-point RSVP LSP to any other PE router in the network as long as it knows that such an LSP is required. To build a full mesh of LSPs between the PE routers requires that each PE router know which of the other PE routers make up the full mesh.

NOTE: In Junos OS, RSVP automatic mesh is configured using the **rsvp-te** configuration statement at the **[edit routing-options dynamic-tunnels dynamic-tunnel-name]** hierarchy level. The **rsvp-te** configuration statement is also available for use in routing instances as a provider-tunnel option. When implemented as a provider-tunnel option, **rsvp-te** is used to configure the RSVP point-to-multipoint LSPs for multiprotocol BGP multicast VPNs, not to configure RSVP automatic mesh.

RELATED DOCUMENTATION

[Configuring RSVP Automatic Mesh | 850](#)

[Example: Configuring RSVP Automatic Mesh | 841](#)

[label-switched-path-template \(Multicast\) | 2859](#)

Path Messages

Each sender host transmits path messages downstream along the routes provided by the unicast and multicast routing protocols. Path messages follow the exact paths of application data, creating path states

in the routers along the way, thus enabling routers to learn the previous-hop and next-hop node for the session. Path messages are sent periodically to refresh path states.

The refresh interval is controlled by a variable called the **refresh-time**, which is the periodical refresh timer expressed in seconds. A path state times out if a router does not receive a specified number of consecutive path messages. This number is specified by a variable called **keep-multiplier**. Path states are kept for $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$ seconds.

Resv Messages

Each receiver host sends reservation request (Resv) messages upstream toward senders and sender applications. Resv messages must follow exactly the reverse path of path messages. Resv messages create and maintain a reservation state in each router along the way.

Resv messages are sent periodically to refresh reservation states. The refresh interval is controlled by the same refresh time variable, and reservation states are kept for $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$ seconds.

PathTear Messages

PathTear messages remove (tear down) path states as well as dependent reservation states in any routers along a path. PathTear messages follow the same path as path messages. A PathTear typically is initiated by a sender application or by a router when its path state times out.

PathTear messages are not required, but they enhance network performance because they release network resources quickly. If PathTear messages are lost or not generated, path states eventually time out when they are not refreshed, and the resources associated with the path are released.

ResvTear Messages

ResvTear messages remove reservation states along a path. These messages travel upstream toward senders of the session. In a sense, ResvTear messages are the reverse of Resv messages. ResvTear messages typically are initiated by a receiver application or by a router when its reservation state times out.

ResvTear messages are not required, but they enhance network performance because they release network resources quickly. If ResvTear messages are lost or not generated, reservation states eventually time out when they are not refreshed, and the resources associated with the reservation are released.

PathErr Messages

When path errors occur (usually because of parameter problems in a path message), the router sends a unicast PathErr message to the sender that issued the path message. PathErr messages are advisory; these messages do not alter any path state along the way.

ResvErr Messages

When a reservation request fails, a ResvErr error message is delivered to all the receivers involved. ResvErr messages are advisory; these messages do not alter any reservation state along the way.

ResvConfirm Messages

Receivers can request confirmation of a reservation request, and this confirmation is sent with a ResvConfirm message. Because of the complex RSVP flow-merging rules, a confirmation message does not necessarily provide end-to-end confirmation of the entire path. Therefore, ResvConfirm messages are an indication, not a guarantee, of potential success.

Juniper Networks routers never request confirmation using the ResvConfirm message; however, a Juniper Networks router can send a ResvConfirm message if it receives a request from another vendor's equipment.

RSVP Reservation Styles

A reservation request includes options for specifying the reservation style. The reservation styles define how reservations for different senders within the same session are treated and how senders are selected.

Two options specify how reservations for different senders within the same session are treated:

- Distinct reservation—Each receiver establishes its own reservation with each upstream sender.
- Shared reservation—All receivers make a single reservation that is shared among many senders.

Two options specify how senders are selected:

- Explicit sender—List all selected senders.
- Wildcard sender—Select all senders, which then participate in the session.

The following reservation styles, formed by a combination of these four options, currently are defined:

- **Fixed filter (FF)**—This reservation style consists of distinct reservations among explicit senders. Examples of applications that use fixed-filter-style reservations are video applications and unicast applications, which both require flows that have a separate reservation for each sender. The fixed filter reservation style is enabled on RSVP LSPs by default.
- **Wildcard filter (WF)**—This reservation style consists of shared reservations among wildcard senders. This type of reservation reserves bandwidth for any and all senders, and propagates upstream toward all senders, automatically extending to new senders as they appear. A sample application for wildcard filter reservations is an audio application in which each sender transmits a distinct data stream. Typically, only a few senders are transmitting at any one time. Such a flow does not require a separate reservation for each sender; a single reservation is sufficient.
- **Shared explicit (SE)**—This reservation style consists of shared reservations among explicit senders. This type of reservation reserves bandwidth for a limited group of senders. A sample application is an audio application similar to that described for wildcard filter reservations.

RSVP Refresh Reduction

RSVP relies on soft-state to maintain the path and reservation states on each router. If the corresponding refresh messages are not sent periodically, the states eventually time out and reservations are deleted. RSVP also sends its control messages as IP datagrams with no reliability guarantee. It relies on periodic refresh messages to handle the occasional loss of Path or Resv messages.

The RSVP refresh reduction extensions, based on RFC 2961, addresses the following problems that result from relying on periodic refresh messages to handle message loss:

- **Scalability**—The scaling problem arises from the periodic transmission and processing overhead of refresh messages, which increases as the number of RSVP sessions increases.
- **Reliability and latency**—The reliability and latency problem stems from the loss of nonrefresh RSVP messages or one-time RSVP messages such as PathTear or PathErr. The time to recover from such a loss is usually tied to refresh interval and the keepalive timer.

The RSVP refresh reduction capability is advertised by enabling the refresh reduction (RR) capable bit in the RSVP common header. This bit is only significant between RSVP neighbors.

RSVP refresh reduction includes the following features:

- RSVP message bundling using the bundle message
- RSVP Message ID to reduce message processing overhead
- Reliable delivery of RSVP messages using Message ID, Message Ack, and Message Nack
- Summary refresh to reduce the amount of information transmitted every refresh interval

The RSVP refresh reduction specification (RFC 2961) allows you to enable some or all of the above capabilities on a router. It also describes various procedures that a router can use to detect the refresh reduction capabilities of its neighbor.

The Junos OS supports all of the refresh reduction extensions, some of which can be selectively enabled or disabled. The Junos OS supports Message ID and therefore can perform reliable message delivery only for Path and Resv messages.

For information about how to configure RSVP refresh reduction, see [“Configuring RSVP Refresh Reduction” on page 829](#).

MTU Signaling in RSVP

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. There are default values for MTUs associated with various protocols. You can also explicitly configure an MTU on an interface.

When an LSP is created across a set of links with different MTU sizes, the ingress router does not know what the smallest MTU is on the LSP path. By default, the MTU for an LSP is 1,500 bytes.

If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router is not aware of this type of traffic loss, because the control plane for the LSP would still function normally.

To prevent this type of packet loss in MPLS LSPs, you can configure MTU signaling in RSVP. This feature is described in RFC 3209. Juniper Networks supports the Integrated Services object for MTU signaling in RSVP. The Integrated Services object is described in RFCs 2210 and 2215. MTU signaling in RSVP is disabled by default.

To avoid packet loss due to MTU mismatches, the ingress router needs to do the following:

- Signal the MTU on the RSVP LSP—To prevent packet loss from an MTU mismatch, the ingress router needs to know what the smallest MTU value is along the path taken by the LSP. Once this MTU value is obtained, the ingress router can assign it to the LSP.
- Fragment packets—Using the assigned MTU value, packets that exceed the size of the MTU can be fragmented into smaller packets on the ingress router before they are encapsulated in MPLS and sent over the RSVP-signaled LSP.

Once both MTU signaling and packet fragmentation have been enabled on an ingress router, any route resolving to an RSVP LSP on this router uses the signaled MTU value. For information about how to configure this feature, see [“Configuring MTU Signaling in RSVP” on page 853](#).

The following sections describe how MTU signaling in RSVP works:

- [How the Correct MTU Is Signaled in RSVP on page 825](#)
- [Determining an Outgoing MTU Value on page 826](#)
- [MTU Signaling in RSVP Limitations on page 826](#)

How the Correct MTU Is Signaled in RSVP

How the correct MTU is signaled in RSVP varies depending on whether the network devices (for example, routers) explicitly support MTU signaling in RSVP or not.

If the network devices support MTU signaling in RSVP, the following occur when you enable MTU signaling:

- The MTU is signaled from the ingress router to the egress router by means of the Adspec object. Before forwarding this object, the ingress router enters the MTU value associated with the interface over which the path message is sent. At each hop in the path, the MTU value in the Adspec object is updated to the minimum of the received value and the value of the outgoing interface.
- The ingress router uses the traffic specification (Tspec) object to specify the parameters for the traffic it is going to send. The MTU value signaled for the Tspec object at the ingress router is the maximum MTU value (9192 bytes for M Series and T Series routers, 9500 bytes for PTX Series Packet Transport Routers). This value does not change en route to the egress router.
- When the Adspec object arrives at the egress router, the MTU value is correct for the path (meaning it is the smallest MTU value discovered). The egress router compares the MTU value in the Adspec object to the MTU value in the Tspec object. It signals the smaller MTU using the Flowspec object in the Resv message.
- When the Resv object arrives at the ingress router, the MTU value in this object is used as the MTU for the next hops that use the LSP.

In a network where there are devices that do not support MTU signaling in RSVP, you might have the following behaviors:

- If the egress router does not support MTU signaling in RSVP, the MTU is set to 1,500 bytes by default.
- A Juniper Networks transit router that does not support MTU signaling in RSVP sets an MTU value of 1,500 bytes in the Adspec object by default.

RELATED DOCUMENTATION

[Media MTU Overview](#)

Determining an Outgoing MTU Value

The outgoing MTU value is the smaller of the values received in the Adspec object compared to the MTU value of the outgoing interface. The MTU value of the outgoing interface is determined as follows:

- If you configure an MTU value under the **[family mpls]** hierarchy level, this value is signaled.
- If you do not configure an MTU, the **inet** MTU is signaled.

MTU Signaling in RSVP Limitations

The following are limitations to MTU signaling in RSVP:

- Changes in the MTU value might cause a temporary loss of traffic in the following situations:
 - For link protection and node protection, the MTU of the bypass is only signaled at the time the bypass becomes active. During the time it takes for the new path MTU to be propagated, packet loss might occur because of an MTU mismatch.
 - For fast reroute, the MTU of the path is updated only after the detour becomes active, causing a delay in an update to the MTU at the ingress router. Until the MTU is updated, packet loss might occur if there is an MTU mismatch.

In both cases, only packets that are larger than the detour or bypass MTU are lost.

- When an MTU is updated, it triggers a change in the next hop. Any change in the next hop causes the route statistics to be lost.
- The minimum MTU supported for MTU signaling in RSVP is 1,488 bytes. This value prevents a false or incorrectly configured value from being used.
- For single-hop LSPs, the MTU value displayed by the **show** commands is the RSVP-signaled value. However, this MPLS value is ignored and the correct IP value is used.

Minimum RSVP Configuration

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

```
rsvp {
  interface interface-name;
}
```

You can include these statements at the following hierarchy levels:

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

To enable RSVP on all interfaces, substitute **all** for the *interface-name* variable.

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

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

You can include this statement at the following hierarchy levels:

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

Configuring RSVP and MPLS

The primary purpose of the Junos RSVP software is to support dynamic signaling within label-switched paths (LSPs). When you enable both MPLS and RSVP on a router, MPLS becomes a client of RSVP. No additional configuration is required to bind MPLS and RSVP.

You can configure MPLS to set up signaled paths by using the **label-switched-path** statement at the [edit protocols mpls] hierarchy level. Each LSP translates into a request for RSVP to initiate an RSVP session. This request is passed through the internal interface between label switching and RSVP. After examining the request information, checking RSVP states, and checking the local routing tables, RSVP initiates one session for each LSP. The session is sourced from the local router and is destined for the target of the LSP.

When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, RSVP notifies MPLS of its status. It is up to MPLS to initiate backup paths or continue retrying the initial path.

To pass label-switching signaling information, RSVP supports four additional objects: Label Request object, Label object, Explicit Route object, and Record Route object. For an LSP to be set up successfully, all routers along the path must support MPLS, RSVP, and the four objects. Of the four objects, the Record Route object is not mandatory.

To configure MPLS and make it a client of RSVP, do the following:

- Enable MPLS on all routers that will participate in the label switching (this is, on all routers that might be part of a label-switching path).
- Enable RSVP on all routers and on all router interfaces that form the LSP.
- Configure the routers at the beginning of the LSP.

Example: Configuring RSVP and MPLS

The following shows a sample configuration for a router at the beginning of an LSP:

```
[edit]
protocols {
  mpls {
    label-switched-path sf-to-london {
      to 192.168.1.4;
    }
  }
  rsvp {
    interface so-0/0/0;
  }
}
```

The following shows a sample configuration for all the other routers that form the LSP:

```
[edit]
protocols {
  mpls {
    interface so-0/0/0;
  }
  rsvp {
    interface so-0/0/0;
  }
}
```


Configuring RSVP Interfaces

IN THIS SECTION

- [Configuring RSVP Refresh Reduction | 829](#)
- [Configuring the RSVP Hello Interval | 831](#)
- [Configuring RSVP Authentication | 832](#)
- [Configuring the Bandwidth Subscription for Class Types | 833](#)
- [Configuring the RSVP Update Threshold on an Interface | 833](#)
- [Configuring RSVP for Unnumbered Interfaces | 834](#)

The following sections describe how to configure RSVP interfaces:

Configuring RSVP Refresh Reduction

You can configure RSVP refresh reduction on each interface by including the following statements in the interface configuration:

- **aggregate** and **reliable**—Enable all RSVP refresh reduction features: RSVP message bundling, RSVP message ID, reliable message delivery, and summary refresh.

In order to have refresh reduction and reliable delivery, you must include the **aggregate** and **reliable** statements.

- **no-aggregate**—Disable RSVP message bundling and summary refresh.
- **no-reliable**—Disable RSVP message ID, reliable message delivery, and summary refresh.

For more information on RSVP refresh reduction, see [“RSVP Refresh Reduction” on page 823](#).

If the **no-reliable** statement is configured on the router (reliable message delivery is disabled), the router accepts RSVP messages that include the Message ID object but ignores the Message ID object and continues performing standard message processing. No error is generated in this case, and RSVP operates normally.

However, not all combinations between two neighbors with different refresh reduction capabilities function correctly. For example, a router is configured with either the **aggregate** statement and **no-reliable** statement or with the **reliable** and **no-aggregate** statements. If an RSVP neighbor sends a Summary Refresh object to this router, no error is generated, but the Summary Refresh object cannot be processed. Consequently, RSVP states can time out on this router if the neighbor is relying only on Summary Refresh to refresh those RSVP states.

We recommend, unless there are specific requirements, that you configure RSVP refresh reduction in a similar manner on each RSVP neighbor.

To enable all RSVP refresh reduction features on an interface, include the **aggregate** statement:

```
aggregate;
```

You can include this statement at the following hierarchy levels:

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

To disable RSVP message bundling and summary refresh, include the **no-aggregate** statement:

```
no-aggregate;
```

You can include this statement at the following hierarchy levels:

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

To enable RSVP message ID and reliable message delivery on an interface, include the **reliable** statement:

```
reliable;
```

You can include this statement at the following hierarchy levels:

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

To disable RSVP message ID, reliable message delivery, and summary refresh, include the **no-reliable** statement:

```
no-reliable;
```

You can include this statement at the following hierarchy levels:

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

Determining the Refresh Reduction Capability of RSVP Neighbors

To determine the RSVP refresh reduction capability of an RSVP neighbor, you need the following information:

- The RR bit advertised by the neighbor
- The local configuration of RSVP refresh reduction
- The actual RSVP messages received from the neighbor

To obtain this information, you can issue a **show rsvp neighbor detail** command. Sample output follows:

```
user@host> show rsvp neighbor detail
```

```
RSVP neighbor: 6 learned
  Address: 192.168.224.178 via: fxp1.0 status: Up
    Last changed time: 10:06, Idle: 5 sec, Up cnt: 1, Down cnt: 0
    Message received: 36
    Hello: sent 69, received: 69, interval: 9 sec
    Remote instance: 0x60b8feba, Local instance: 0x74bc7a8d
    Refresh reduction: not operational

  Address: 192.168.224.186 via: fxp2.0 status: Down
    Last changed time: 10:17, Idle: 40 sec, Up cnt: 0, Down cnt: 0
    Message received: 6
    Hello: sent 20, received: 0, interval: 9 sec
    Remote instance: 0x0, Local instance: 0x2aeb339
    Refresh reduction: incomplete
      Remote end: disabled, Ack-extension: enabled

  Address: 192.168.224.188 via: fxp2.0 status: Up
    Last changed time: 4:15, Idle: 0 sec, Up cnt: 1, Down cnt: 0
    Message received: 55
    Hello: sent 47, received: 31, interval: 9 sec
    Remote instance: 0x6436a35b, Local instance: 0x663849f0
    Refresh reduction: operational
      Remote end: enabled, Ack-extension: enabled
```

For more information on the [show rsvp neighbor detail](#) command.

Configuring the RSVP Hello Interval

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello

packets are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

For Juniper Networks routers, configuring a shorter or longer RSVP hello interval has no impact on whether or not an RSVP session is brought down. RSVP sessions are kept up even if RSVP hello packets are no longer being received. RSVP sessions are maintained until either the router stops receiving IGP hello packets or the RSVP Path and Resv messages time out. However, starting from Junos OS Release 16.1, when RSVP hello messages time-out, the RSVP sessions are brought down.

The RSVP hello interval might also be impacted when another vendor's equipment brings down an RSVP session. For example, a neighboring non-Juniper Networks router might be configured to monitor RSVP hello packets.

To modify how often RSVP sends hello packets, include the **hello-interval** statement:

```
hello-interval seconds;
```

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

Configuring RSVP Authentication

All RSVP protocol exchanges can be authenticated to guarantee that only trusted neighbors participate in setting up reservations. By default, RSVP authentication is disabled.

RSVP authentication uses a Hashed Message Authentication Code (HMAC)-MD5 message-based digest. This scheme produces a message digest based on a secret authentication key and the message contents. (The message contents also include a sequence number.) The computed digest is transmitted with RSVP messages. Once you have configured authentication, all received and transmitted RSVP messages with all neighbors are authenticated on this interface.

MD5 authentication provides protection against forgery and message modification. It also can prevent replay attacks. However, it does not provide confidentiality, because all messages are sent in clear text.

By default, authentication is disabled. To enable authentication, configure a key on each interface by including the **authentication-key** statement:

```
authentication-key key;
```

You can include this statement at the following hierarchy levels:

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

Configuring the Bandwidth Subscription for Class Types

By default, RSVP allows 100 percent of the bandwidth for a class type to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

For detailed instructions on how to configure the bandwidth subscription for class types, see [“Configuring the Bandwidth Subscription Percentage for LSPs” on page 1210](#).

Configuring the RSVP Update Threshold on an Interface

The interior gateway protocols (IGPs) maintain the traffic engineering database, but the current available bandwidth on the traffic engineering database links originates from RSVP. When a link's bandwidth changes, RSVP informs the IGPs, which can then update the traffic engineering database and forward the new bandwidth information to all network nodes. The network nodes then know how much bandwidth is available on the traffic engineering database link (local or remote), and CSPF can correctly compute the paths.

However, IGP updates can consume excessive system resources. Depending on the number of nodes in a network, it might not be desirable to perform an IGP update for small changes in bandwidth. By configuring the **update-threshold** statement at the **[edit protocols rsvp]** hierarchy level, you can adjust the threshold at which a change in the reserved bandwidth triggers an IGP update.

You can configure a value of from 0.001 percent through 20 percent (the default is 10 percent) for when to trigger an IGP update. If the change in the reserved bandwidth is greater than or equal to the configured threshold percentage of the static bandwidth on that interface, then an IGP update occurs. For example, if you have configured the **update-threshold** statement to be 15 percent and the router discovers that the reserved bandwidth on a link has changed by 10 percent of the link bandwidth, RSVP does not trigger an IGP update. However, if the reserved bandwidth on a link changes by 20 percent of the link bandwidth, RSVP triggers an IGP update.

You can also configure the threshold as an absolute value using the **threshold-value** option under the **update-threshold** statement.

If the threshold-value is configured to greater than 20% of bandwidth on that link, the threshold-value is capped at 20% of bandwidth.

For instance, if bandwidth on an interface is 1Gbps, and the **threshold-value** is configured greater than 200Mbps, the **threshold-value** is capped at 200Mbps. The *threshold-percent* is displayed as 20.000% and the **threshold-value** as 200Mbps.

NOTE: The two options, *threshold-percent* and **threshold-value**, are mutually exclusive. You can configure only one option at a given point in time to generate an IGP update for lower bandwidth reservations. As a result, when one option is configured, the other option is calculated and displayed on the CLI.

For instance, on a link of 1Gbps, if the *threshold-percent* is configured to 5%, the **threshold-value** is calculated and displayed as 50Mbps. Similarly, if the **threshold-value** is configured to 50m, then the *threshold-percent* is calculated and displayed as 5%.

To adjust the threshold at which a change in the reserved bandwidth triggers an IGP update, include the **update-threshold** statement. Because of the update threshold, it is possible for Constrained Shortest Path First (CSPF) to compute a path using outdated traffic engineering database bandwidth information on a link. If RSVP attempts to establish an LSP over that path, it might find that there is insufficient bandwidth on that link. When this happens, RSVP triggers an IGP traffic engineering database update, flooding the updated bandwidth information on the network. CSPF can then recompute the path by using the updated bandwidth information, and attempt to find a different path, avoiding the congested link. Note that this functionality is the default and does not need any additional configuration.

You can configure the **rsvp-error-hold-time** statement at the **[edit protocols mpls]** hierarchy level or the **[edit logical-systems logical-system-name protocols mpls]** hierarchy level to improve the accuracy of the traffic engineering database (including the accuracy of bandwidth estimates for LSPs) using information provided by PathErr messages. See [“Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages” on page 1189](#).

Configuring RSVP for Unnumbered Interfaces

The Junos OS supports RSVP traffic engineering over unnumbered interfaces. Traffic engineering information about unnumbered links is carried in the IGP traffic engineering extensions for OSPF and IS-IS as described in RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*, and RFC 4205, *Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*. Unnumbered links can also be specified in the MPLS traffic engineering signaling as described in RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*. This feature allows you avoid having to configure IP addresses for each interface participating in the RSVP-signaled network.

To configure RSVP for unnumbered interfaces, you must configure the router with a router ID using the **router-id** statement specified at the **[edit routing-options]** hierarchy level. The router ID must be available for routing (you can typically use the loopback address). The RSVP control messages for the unnumbered links are sent using the router ID address (rather than a randomly selected address).

To configure link protection and fast reroute on a router with unnumbered interfaces enabled, you must configure at least two addresses. We recommend that you configure a secondary interface on the loopback in addition to configuring the router ID.

Release History Table

Release	Description
19.4R1	
16.1	However, starting from Junos OS Release 16.1, when RSVP hello messages time-out, the RSVP sessions are brought down.

Configuring RSVP Node-ID Hellos

You can configure node-ID based RSVP hellos to ensure that Juniper Networks routers can interoperate with the equipment of other vendors. By default, Junos OS uses interface-based RSVP hellos. Node-ID based RSVP hellos are specified in RFC 4558, *Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement*. RSVP node-ID hellos are useful if you have configured BFD to detect problems over RSVP interfaces, allowing you to disable interface hellos for these interfaces. You can also use node-ID hellos for graceful restart procedures.

Node-ID hellos can be enabled globally for all RSVP neighbors. By default, node-ID hello support is disabled. If you have not enabled RSVP node IDs on the router, the Junos OS does not accept any node-ID hello packets.

To enable RSVP node-ID hellos globally on the router, include the [node-hello](#) statement at the following hierarchy levels:

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

You can also explicitly disable RSVP interface hellos globally. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the [hello-interval](#) statement. This configuration disables RSVP interface hellos globally, but enables RSVP interface hellos on the specified interface (the RSVP interface you configure the [hello-interval](#) statement on). This configuration might be necessary in a heterogeneous network in which some devices support RSVP node-ID hellos and other devices support RSVP interface hellos.

To disable RSVP interface hellos globally on the router, include the [no-interface-hello](#) statement at the following hierarchy levels:

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

RELATED DOCUMENTATION

[no-interface-hello](#) | 2872

[node-hello](#) | 2867

Example: Configuring RSVP-Signaled LSPs

IN THIS SECTION

- [Requirements](#) | 836
- [Overview and Topology](#) | 836
- [Configuration](#) | 837
- [Verification](#) | 839

This example shows how to create an LSP between routers in an IP network using RSVP as the signaling protocol.


Requirements

Before you begin, delete security services from the device. See *Example: Deleting Security Services*.

Overview and Topology

Using RSVP as a signaling protocol, you can create LSPs between routers in an IP network. In this example, you configure a sample network as shown in [Figure 61 on page 837](#).

Figure 61: Typical RSVP-Signaled LSP



To establish an LSP between routers, you must individually enable the MPLS family and configure RSVP on each of the transit interfaces in the MPLS network. This example shows how to enable MPLS and configure RSVP on the ge-0/0/0 transit interface. Additionally, you must enable the MPLS process on all of the MPLS interfaces in the network.

This example shows how to define an LSP from R1 to R7 on the ingress router (R1) using R7's loopback address (10.0.9.7). The configuration reserves 10 Mbps of bandwidth. Additionally, the configuration disables the CSPF algorithm, ensuring that Hosts C1 and C2 use the RSVP-signaled LSP that correspond to the network IGP's shortest path.

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.

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols rsvp interface ge-0/0/0.0
set protocols mpls label-switched-path r1-r7 to 10.0.9.7
set protocols mpls label-switched-path r1-r7 bandwidth 10m
set protocols mpls interface all
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure RSVP:

1. Enable the MPLS family on all transit interfaces in the MPLS network.

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family mpls
```

2. Enable RSVP on each transit interface in the MPLS network.

```
[edit]
user @host# set protocols rsvp interface ge-0/0/0
```

3. Enable the MPLS process on the transit interface in the MPLS network.

```
[edit]
user@host# set protocols mpls interface ge-0/0/0
```

4. Define the LSP on the ingress router.

```
[edit protocols mpls]
user@host# set label-switched-path r1-r7 to 10.0.9.7
```

5. Reserve 10 Mbps of bandwidth on the LSP.

```
[edit protocols mpls]
user @host# set label-switched-path r1-r7 bandwidth 10m
```

Results

Confirm your configuration by entering the **show** command from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
user@host# show
...
```



```

interfaces {
  ge-0/0/0 {
    family mpls;
  }
}
...
protocols {
  rsvp {
    interface ge-0/0/0.0;
  }
  mpls {
    label-switched-path r1-r7 {
      to 10.0.9.7;
      bandwidth 10m;
    }
    interface all;
  }
}
...

```

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

Verification

IN THIS SECTION

- [Verifying RSVP Neighbors | 839](#)
- [Verifying RSVP Sessions | 840](#)
- [Verifying the Presence of RSVP-Signaled LSPs | 841](#)

To confirm that the configuration is working properly, perform these tasks:

Verifying RSVP Neighbors

Purpose

Verify that each device shows the appropriate RSVP neighbors—for example, that Router R1 in [Figure 61 on page 837](#) lists both Router R3 and Router R2 as RSVP neighbors.

Action

From the CLI, enter the **show rsvp neighbor** command.

```
user@r1> show rsvp neighbor
```

```

RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx
10.0.6.2         0   3/2      13:01         3    366/349
10.0.3.3         0   1/0      22:49         3    448/448

```

The output shows the IP addresses of the neighboring routers. Verify that each neighboring RSVP router loopback address is listed.

Verifying RSVP Sessions

Purpose

Verify that an RSVP session has been established between all RSVP neighbors. Also, verify that the bandwidth reservation value is active.

Action

From the CLI, enter the **show rsvp session detail** command.

```
user@r1> show rsvp session detail
```

```

Ingress RSVP: 1 sessions

10.0.9.7
  From: 10.0.6.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-r7, LSPpath: Primary
  Bidirectional, Upstream label in: -, Upstream label out: -
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100000
  Resv style: 1 FF, Label in: -, Label out: 100000
  Time left:    -, Since: Thu Jan 26 17:57:45 2002
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 3 receiver 17 protocol 0
  PATH rcvfrom: localclient
  PATH sentto: 10.0.4.13 (ge-0/0/1.0) 1467 pkts
  RESV rcvfrom: 10.0.4.13 (ge-0/0/1.0) 1467 pkts
  Record route: <self> 10.0.4.13 10.0.2.1 10.0.8.10

```

The output shows the detailed information, including session IDs, bandwidth reservation, and next-hop addresses, for each established RSVP session. Verify the following information:

- Each RSVP neighbor address has an entry for each neighbor, listed by loopback address.
- The state for each LSP session is **Up**.
- For **Tspec**, the appropriate bandwidth value, **10Mbps**, appears.

Verifying the Presence of RSVP-Signaled LSPs

Purpose

Verify that the routing table of the entry (ingress) router has a configured LSP to the loopback address of the other router. For example, verify that the **inet.3** routing table of the R1 entry router in [Figure 61 on page 837](#) has a configured LSP to the loopback address of Router R7.

Action

From the CLI, enter the **show route table inet.3** command.

```
user@r1> show route table inet.3
```

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.9.7/32          *[RSVP/7] 00:05:29, metric 10
                    > to 10.0.4.17 via ge-0/0/0.0, label-switched-path r1-r7
```

The output shows the RSVP routes that exist in the **inet.3** routing table. Verify that an RSVP-signaled LSP is associated with the loopback address of the exit (egress) router, R7, in the MPLS network.

RELATED DOCUMENTATION

[Understanding the RSVP Signaling Protocol | 810](#)

[CLI Explorer](#)

Example: Configuring RSVP Automatic Mesh

IN THIS SECTION

- [Requirements | 842](#)
- [Overview | 842](#)

●	Configuration 843
●	Verification 845

Service providers often use BGP and MPLS VPNs to efficiently scale the network while delivering revenue-generating services. In these environments, BGP is used to distribute the VPN routing information across the service provider's network, while MPLS is used to forward that VPN traffic from one VPN site to another.

When adding a new PE router that will participate in BGP and MPLS VPNs, all of the previously existing PE routers must be configured to peer with the new PE router for both BGP and MPLS. As each new PE router is added to the service provider's network, the configuration burden soon becomes too much to handle.

The configuration requirements for BGP peering can be reduced with the use of route reflectors. In RSVP signaled MPLS networks, RSVP automatic mesh can minimize the configuration burden for the MPLS portion of the network. Configuring **rsvp-te** on all PE routers allows RSVP to automatically create the needed LSPs when a new PE router is added.

Requirements

This example uses the following hardware and software components:

- A router running Junos OS Release 10.1 or later.
- A BGP and MPLS VPN using RSVP as the MPLS label-switched path (LSP) signaling protocol.

Overview


This example shows how to configure RSVP automatic mesh on a PE router using the **rsvp-te** configuration statement. In order for RSVP automatic mesh to function properly, all of the PE routers in the full mesh configuration must have the **rsvp-te** statement configured. This ensures that any new PE routers that are added later will also benefit from the automatic mesh feature, provided that they too are configured with the **rsvp-te** statement.

Given this requirement, this example only shows the configuration on the newly added PE router. It is assumed that RSVP automatic mesh has already been configured on the existing PE routers.

Topology

In [Figure 62 on page 843](#), there are three existing PE routers, PE1, PE2, and PE3, in the topology. PE4 has been added, and RSVP automatic mesh will be configured. The cloud represents the service provider network, and the network address, 192.0.2.0/24, is shown in the center of the figure.

Figure 62: Service Provider Network with PE Routers



Configuration

IN THIS SECTION

- [Configuring RSVP Automatic Mesh | 844](#)
- [Results | 844](#)

Configuring RSVP automatic mesh involves performing these tasks:

- Enabling the **rsvp-te** configuration statement at the **[edit routing-options dynamic-tunnels dynamic-tunnel-name]** hierarchy level.
- Configuring the required **destination-networks** element.

This configuration element specifies the IPv4 prefix range for the destination network. Only tunnels within the specified prefix range can be created.

- Configuring the required **label-switched-path-template** element.

This configuration element takes either **default-template** or the name of a preconfigured LSP template as an argument. The **default-template** is a system-defined template that requires no user 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.

PE4 Router


```
set routing-options dynamic-tunnels dt-1 rsvp-te rsvp-te-1 label-switched-path-template
  default-template
set routing-options dynamic-tunnels dt-1 rsvp-te rsvp-te-1 destination-networks 192.0.2.0/24
```

Configuring RSVP Automatic Mesh

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To enable RSVP automatic mesh:

1. Configure **rsvp-te** at the **[edit routing-options dynamic-tunnels]** hierarchy level.

```
[edit routing-options dynamic-tunnels]
user@PE4# set dt-1 rsvp-te rsvp-te-1 label-switched-path-template default-template
```

2. Configure **destination-networks** at the **[edit routing-options dynamic-tunnels]** hierarchy level.

```
[edit routing-options dynamic-tunnels]
user@PE4# set dt-1 rsvp-te rsvp-te-1 destination-networks 192.0.2.0/24
```

Results

Issue the **show** command from the **[edit routing-options dynamic-tunnels]** hierarchy level to see the results of your configuration:

```
[edit routing-options dynamic-tunnels]
user@PE4#show
dt-1 {
  rsvp-te rsvp-te-1 {
    label-switched-path-template {
      default-template;
    }
    destination-networks {
      192.0.2.0/24;
    }
  }
}
```


Verification

IN THIS SECTION

- [Verifying the Existence of RSVP Automatic Mesh Tunnels on Router PE4 | 845](#)
- [Verifying the Existence of MPLS LSPs on Router PE4 | 845](#)

Verifying the Existence of RSVP Automatic Mesh Tunnels on Router PE4

Purpose

To verify the operation of the newly configured PE4 router, issue the **show dynamic-tunnels database** command from operational mode. This command will show the destination network to which dynamic tunnels can be created.

Action

```
user@PE4> show dynamic-tunnels database
Table: inet.3
Destination-network: 192.0.2.0/24
```

Verifying the Existence of MPLS LSPs on Router PE4

Purpose

To verify the existence of MPLS LSPs on the PE4 router, issue the **show mpls lsp** command from operational mode. This command will show the state of the MPLS LSPs.

Action

```
user@PE4> show mpls lsp
```

```
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Egress LSP: 3 sessions
To From State Rt Style Labelin Labelout LSPname
192.0.2.104 192.0.2.103 Up 0 1 FF 3 - PE2-PE4
192.0.2.104 192.0.2.102 Up 0 1 FF 3 - PE2-PE4
192.0.2.104 192.0.2.101 Up 0 1 FF 3 - PE1-PE4
Total 3 displayed, Up 3, Down 0
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```


RELATED DOCUMENTATION

[Understanding RSVP Automatic Mesh | 819](#)

[rsvp-te | 2888](#)

Configuring Hello Acknowledgments for Nonsession RSVP Neighbors

The **hello-acknowledgements** statement controls the hello acknowledgment behavior between RSVP neighbors regardless of whether or not they are in the same session.

Hello messages received from RSVP neighbors that are not part of a common RSVP session are discarded. If you configure the **hello-acknowledgements** statement at the **[edit protocols rsvp]** hierarchy level, hello messages from nonsession neighbors are acknowledged with a hello acknowledgment message. When hellos are received from nonsession neighbors, an RSVP neighbor relationship is created and periodic hello messages can now be received from the nonsession neighbor. The **hello-acknowledgements** statement is disabled by default. Configuring this statement allows RSVP-capable routers to be discovered using hello packets and verifies that the interface is able to receive RSVP packets before sending any MPLS LSP setup messages.

Once you enable hello acknowledgments for nonsession RSVP neighbors, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or you change the configuration. Interface-based neighbors are not automatically aged out.

```
hello-acknowledgements;
```

You can include this statement at the following hierarchy levels:

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

Switching LSPs Away from a Network Node

You can configure the router to switch active LSPs away from a network node using a bypass LSP enabled for an interface. This feature might be used to maintain active networks when a device needs to be replaced without interrupting traffic transiting the network. The LSPs can be either static or dynamic.

1. You first need to configure either link or node protection for the traffic that needs to pass around the network device you intend to disable. To function properly, the bypass LSP must use a different logical interface than the protected LSP.

2. To prepare the router to begin switching traffic away from a network node, configure the **always-mark-connection-protection-tlv** statement:

```
always-mark-connection-protection-tlv;
```

The router then marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality.

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls interface *interface-name*]
 - [edit logical-systems *logical-system-name* protocols mpls interface *interface-name*]
3. You then need to configure the **switch-away-lsps** statement to switch the traffic from the protected LSP to the bypass LSP, effectively bypassing the default downstream network device. The actual link itself is not brought down by this configuration.

To configure the router to switch traffic away from a network node, configure the **switch-away-lsps** statement:

```
switch-away-lsps;
```

You can configure this statement at the following hierarchy levels:

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

Note the following limitations related to switching active LSPs away from a network node:

- The switch-away feature is supported on MX Series routers only.
- The switch-away feature is not supported for switching traffic from primary point-to-multipoint LSPs to bypass point-to-multipoint LSPs. If you configure the **switch-away-lsps** statement for a point-to-multipoint LSP, traffic is not switched to the bypass point-to-multipoint LSP.
- If you configure the switch-away feature on an interface along the path of a dynamic LSP, new dynamic LSPs cannot be established over that path. The switch-away feature prevents the make-before-break behavior of RSVP-signaled LSPs. The make-before-break behavior normally causes the router to first attempt to re-signal a dynamic LSP before tearing down the original.

RELATED DOCUMENTATION

[Configuring Node Protection or Link Protection for LSPs](#) | 879

Configuring RSVP Setup Protection

You can configure the facility-backup fast reroute mechanism to provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. This feature is applicable in the following scenario:

1. A failed link or node is present on the strict explicit path of an LSP before the LSP is signaled.
2. There is also a bypass LSP protecting the link or node.
3. RSVP signals the LSP through the bypass LSP. The LSP appears as if it was originally set up along its primary path and then failed over to the bypass LSP because of the link or node failure.
4. When the link or node has recovered, the LSP can be automatically reverted to the primary path.

You should configure the **setup-protection** statement at the **[edit protocols rsvp]** on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path. You can issue a **show rsvp session** command to determine whether or not the LSP has setup protection enabled on a router acting as a point of local repair (PLR) or a merge point.

To enable RSVP setup protection, include the **setup-protection** statement

```
setup-protection;
```

You can include this statement at the following hierarchy levels:

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

Configuring Load Balancing Across RSVP LSPs

By default, when you have configured several RSVP LSPs to the same egress router, the LSP with the lowest metric is selected and carries all traffic. If all of the LSPs have the same metric, one of the LSPs is selected at random and all traffic is forwarded over it.

Alternatively, you can load-balance traffic across all of the LSPs by enabling per-packet load balancing.

To enable per-packet load balancing on an ingress LSP, configure the **policy-statement** statement as follows:

```
[edit policy-options]
policy-statement policy-name {
```



```

    then {
        load-balance per-packet;
    }
    accept;
}

```

You then need to apply this statement as an export policy to the forwarding table.

Once per-packet load balancing is applied, traffic is distributed equally between the LSPs (by default).

You need to configure per-packet load balancing if you want to enable PFE fast reroute. To enable PFE fast reroute, include the **policy-statement** statement for per-packet load balancing shown in this section in the configuration of each of the routers where a reroute might take place. See also [“Configuring Fast Reroute” on page 479](#).

You can also load-balance the traffic between the LSPs in proportion to the amount of bandwidth configured for each LSP. This capability can better distribute traffic in networks with asymmetric bandwidth capabilities across external links, since the configured bandwidth of an LSP typically reflects the traffic capacity of that LSP.

To configure RSVP LSP load balancing, include the **load-balance** statement with the **bandwidth** option:

```

load-balance {
    bandwidth;
}

```

You can configure this statement at the following hierarchy levels:

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

Keep the following information in mind when you use the **load-balance** statement:

- If you configure the **load-balance** statement, the behavior of currently running LSPs is not altered. To force currently running LSPs to use the new behavior, you can issue a **clear mpls lsp** command.
- The **load-balance** statement only applies to ingress LSPs that have per-packet load balancing enabled.
- For Differentiated Services-aware traffic engineered LSPs, the bandwidth of an LSP is calculated by summing the bandwidth of all of the class types.

Configuring RSVP Automatic Mesh

You can configure RSVP to establish point-to-point label-switched paths (LSPs) automatically for any new PE router added to a full mesh of LSPs. To enable this feature, you must configure the **rsvp-te** statement on all of the PE routers in the full mesh.

NOTE: You cannot configure RSVP automatic mesh in conjunction with CCC. CCC cannot use the dynamically generated LSPs.

To configure RSVP automatic mesh, include the **rsvp-te** statement:

```
rsvp-te {
  destination-networks network-prefix;
  label-switched-path-template (Multicast) {
    default-template;
    template-name;
  }
}
```

You can configure these statements at the following hierarchy levels:

- [edit routing-options dynamic-tunnels *tunnel-name*]
- [edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*]

You must also configure the following statements to enable RSVP automatic mesh:

- **destination-networks**—Specify the IP version 4 (IPv4) prefix range for the destination network. Dynamic tunnels within the specified IPv4 prefix range can be initiated.
- **label-switched-path-template (Multicast)**—You can configure either the default template explicitly using the **default-template** option, or you can configure an LSP template of your own using the **template-name** option. The LSP template acts as a model configuration for the dynamically generated LSPs.

RELATED DOCUMENTATION

Configuring Dynamic Point-to-Multipoint Flooding LSPs

[Example: Configuring RSVP Automatic Mesh](#) | 841

Configuring Timers for RSVP Refresh Messages

RSVP uses two related timing parameters:

- **refresh-time**—The refresh time controls the interval between the generation of successive refresh messages. The default value for the refresh time is 45 seconds. This number is derived from the **refresh-time** statement's default value of 30, multiplied by a fixed value of 1.5. This computation differs from RFC 2205, which states that the refresh time should be multiplied by a random value in the range from 0.5 through 1.5.

Refresh messages include path and Resv messages. Refresh messages are sent periodically so that reservation states in neighboring nodes do not time out. Each path and Resv message carries the refresh timer value, and the receiving node extracts this value from the messages.

- **keep-multiplier**—The keep multiplier is a small, locally configured integer from 1 through 255. The default value is 3. It indicates the number of messages that can be lost before a particular state is declared stale and must be deleted. The keep multiplier directly affects the lifetime of an RSVP state.

To determine the lifetime of a reservation state, use the following formula:

```
lifetime = (keep-multiplier + 0.5) x (1.5 x refresh-time)
```

In the worst case, (**keep-multiplier** – 1) successive refresh messages must be lost before a reservation state is deleted.

We do not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

By default, the refresh timer value is 30 seconds. To modify this value, include the **refresh-time** statement:

```
refresh-time seconds;
```

You can include this statement at the following hierarchy levels:

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

The default value of the keep multiplier is 3. To modify this value, include the **keep-multiplier** statement:

```
keep-multiplier number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]

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

Preempting RSVP Sessions

Whenever bandwidth is insufficient to handle all RSVP sessions, you can control the preemption of RSVP sessions. By default, an RSVP session is preempted only by a new higher-priority session.

To always preempt a session when the bandwidth is insufficient, include the **preemption** statement with the **aggressive** option:

```
preemption aggressive;
```

You can include this statement at the following hierarchy levels:

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

To disable RSVP session preemption, include the **preemption** statement with the **disabled** option:

```
preemption disabled;
```

To return to the default (that is, preempt a session only for a new higher-priority session), include the **preemption** statement with the **normal** option:

```
preemption normal;
```

You can include this statement at the following hierarchy levels:

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

Configuring MTU Signaling in RSVP

IN THIS SECTION

- [Enabling MTU Signaling in RSVP | 853](#)
- [Enabling Packet Fragmentation | 854](#)

To configure maximum transmission unit (MTU) signaling in RSVP, you need to configure MPLS to allow IP packets to be fragmented before they are encapsulated in MPLS. You also need to configure MTU signaling in RSVP. For troubleshooting purposes, you can configure MTU signaling alone without enabling packet fragmentation.

To configure MTU signaling in RSVP, include the **path-mtu** statement:

```
path-mtu {  
    allow-fragmentation;  
    rsvp {  
        mtu-signaling;  
    }  
}
```

You can include this statement at the following hierarchy levels:

- [\[edit protocols mpls\]](#)
- [\[edit logical-systems *logical-system-name* protocols mpls\]](#)

The following sections describe how to enable packet fragmentation and MTU signaling in RSVP:

Enabling MTU Signaling in RSVP

To enable MTU signaling in RSVP, include the **rsvp mtu-signaling** statement:

```
rsvp mtu-signaling;
```

You can include this statement at the following hierarchy levels:

- [\[edit protocols mpls path-mtu\]](#)
- [\[edit logical-systems *logical-system-name* protocols mpls path-mtu\]](#)

Once you have committed the configuration, changes in the MTU signaling behavior for RSVP take effect the next time the path is refreshed.

You can configure the **mtu-signaling** statement by itself at the **[edit protocols mpls path-mtu rsvp]** hierarchy level. This can be useful for troubleshooting. If you configure just the **mtu-signaling** statement, you can use the **show rsvp session detail** command to determine what the smallest MTU is on an LSP. The **show rsvp session detail** command displays the MTU value received and sent in the Adspec object.

Enabling Packet Fragmentation

To allow IP packets to be fragmented before they are encapsulated in MPLS, include the **allow-fragmentation** statement:

```
allow-fragmentation;
```

You can include this statement at the following hierarchy levels:


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

NOTE: Do not configure the **allow-fragmentation** statement alone. Always configure it in conjunction with the **mtu-signaling** statement.

Configuring Ultimate-Hop Popping for LSPs


By default, RSVP-signaled LSPs use penultimate-hop popping (PHP). [Figure 47 on page 600](#) illustrates a penultimate-hop popping LSP between Router PE1 and Router PE2. Router CE1 forwards a packet to its next hop (Router PE1), which is also the LSP ingress. Router PE1 pushes label 1 on the packet and forwards the labeled packet to Router P1. Router P1 completes the standard MPLS label swapping operation, swapping label 1 for label 2, and forwards the packet to Router P2. Since Router P2 is the penultimate-hop router for the LSP to Router PE2, it first pops the label and then forwards the packet to Router PE2. When Router PE2 receives it, the packet can have a service label, an explicit-null label, or just be a plain IP or VPLS packet. Router PE2 forwards the unlabeled packet to Router CE2.

Figure 63: Penultimate-Hop Popping for an LSP



You can also configure ultimate-hop popping (UHP) (as shown in [Figure 48 on page 600](#)) for RSVP-signaled LSPs. Some network applications can require that packets arrive at the egress router (Router PE2) with a non-null outer label. For an ultimate-hop popping LSP, the penultimate router (Router P2 in [Figure 48 on page 600](#)) performs the standard MPLS label swapping operation (in this example, label 2 for label 3) before forwarding the packet to egress Router PE2. Router PE2 pops the outer label and performs a second lookup of the packet address to determine the end destination. It then forwards the packet to the appropriate destination (either Router CE2 or Router CE4).

Figure 64: Ultimate-Hop Popping for an LSP



The following network applications require that you configure UHP LSPs:

- MPLS-TP for performance monitoring and in-band OAM
- Edge protection virtual circuits

The following features do not support the UHP behavior:

- LDP-signaled LSPs
- Static LSPs
- Point-to-multipoint LSPs
- CCC
- **traceroute** command

For more information about UHP behavior, see Internet draft [draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt](#), *Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs*.

For point-to-point RSVP-signaled LSPs, UHP behavior is signaled from the LSP ingress. Based on the ingress router configuration, RSVP can signal the UHP LSP with the non-PHP flag set. RSVP PATH messages carry the two flags in the LSP-ATTRIBUTES object. When the egress router receives the PATH message, it assigns a non-null label to the LSP. RSVP also creates and installs two routes in the mpls.0 routing table. S refers to the S bit of the MPLS label, which indicates whether or not the bottom of the label stack has been reached.

- Route S=0—Indicates that there are more labels in the stack. The next hop for this route points to the mpls.0 routing table, triggering a chained MPLS label lookup to discover the remaining MPLS labels in the stack.
- Route S=1—Indicates that there are no more labels. The next hop points to the inet.0 routing table if the platform supports chained and multi-family lookup. Alternatively, the label route can point to a VT interface to initiate IP forwarding.

If you enable UHP LSPs, MPLS applications such as Layer 3 VPNs, VPLS, Layer 2 VPNs, and Layer 2 circuits can use the UHP LSPs. The following explains how UHP LSPs affect the different types of MPLS applications:

- Layer 2 VPNs and Layer 2 circuits—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label (S=1) is the VC label. A lookup based on the transport label results in a table handle for the mpls.0 routing table. There is an additional route in the mpls.0 routing table corresponding to the inner label. A lookup based on the inner label results in the CE router next hop.
- Layer 3 VPN—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label is the VPN label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. There are two cases in this scenario. By default, Layer 3 VPNs advertise the per-next hop label. A lookup based on the inner label results in the next hop toward the CE router. However, if you have configured the **vrf-table-label** statement for the Layer 3 VPN routing instance, the inner LSI label points to the VRF routing table. An IP lookup is also completed for the VRF routing table.

NOTE: UHP for Layer 3 VPNs configured with the **vrf-table-label** statement is supported on MX Series 5G Universal Routing Platforms only.

- VPLS—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0) and the inner label is the VPLS label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. A lookup based on the inner label in mpls.0 routing table results in the LSI tunnel interface of the VPLS routing instance if tunnel-services is not configured (or a VT interface not available). MX 3D Series routers support chained lookup and multi-family lookup.

NOTE: UHP for VPLS configured with the **no-tunnel-service** statement is supported on MX 3D Series routers only.

- **IPv4 over MPLS**—A packet arrives at the PE router (egress of the UHP LSP) with one label ($S=1$). A lookup based on this label returns a VT tunnel interface. Another IP lookup is completed on the VT interface to determine where to forward the packet. If the routing platform supports multi-family and chained lookups (for example, MX 3D routers and PTX Series Packet Transport Routers), lookup based on label route ($S=1$) points to the inet.0 routing table.
- **IPv6 over MPLS**—For IPv6 tunneling over MPLS, PE routers advertise IPv6 routes to each other with a label value of 2. This is the explicit null label for IPv6. As a result, the forwarding next hops for IPv6 routes that are learned from remote PE routers normally push two labels. The inner label is 2 (it could be different if the advertising PE router is from another vendor), and the router label is the LSP label. Packets arrive at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label ($S=0$), and the inner label is the IPv6 explicit-null label (label 2). Lookup based on the inner label in the mpls.0 routing table redirects back to the mpls.0 routing table. On MX 3D Series routers, the inner label (label 2) is stripped off and an IPv6 lookup is done using the inet6.0 routing table.
- **Enabling both PHP and UHP LSPs**—You can configure both PHP and UHP LSPs over the same network paths. You can separate PHP and UHP traffic by selecting forwarding LSP next hops using a regular expression with the **install-nexthop** statement. You can also separate traffic by simply naming the LSPs appropriately.

The following statements enable ultimate-hop popping for an LSP. You can enable this feature on a specific LSP or for all of the ingress LSPs configured on the router. Configure these statements on the router at the LSP ingress.

1. To enable ultimate-hop popping, include the **ultimate-hop-popping** statement:

```
ultimate-hop-popping;
```

Include this statement at the **[edit protocols mpls label-switched-path label-switched-path-name]** hierarchy level to enable ultimate-hop popping on a specific LSP. Include this statement at the **[edit protocols mpls]** hierarchy level to enable ultimate-hop popping on all of the ingress LSPs configured on the router. You can also configure the **ultimate-hop-popping** statement under the equivalent **[edit logical-routers]** hierarchy levels.

NOTE: When you enable ultimate-hop popping, RSVP attempts to resignal existing LSPs as ultimate-hop popping LSPs in a make-before-break fashion. If an egress router does not support ultimate-hop popping, the existing LSP is torn down (RSVP sends a PathTear message along an LSP's path, removing the path state and dependent reservation state and releasing the associated networking resources).

If you disable ultimate-hop popping, RSVP resignals existing LSPs as penultimate-hop popping LSPs in a make-before-break fashion.

2. If you want to enable both ultimate-hop-popping and chained next hops on MX 3D Series routers only, you also need to configure the **enhanced-ip** option for the **network-services** statement:

```
network-services enhanced-ip;
```

You configure this statement at the **[edit chassis]** hierarchy level. Once you have configured the **network-services** statement, you need to reboot the router to enable UHP behavior.

RELATED DOCUMENTATION

[MPLS Label Allocation | 420](#)

[Configuring Corouted Bidirectional LSPs | 570](#)

[network-services](#)

[ultimate-hop-popping | 2824](#)

Configuring RSVP to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of an LSP. The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. When ultimate-hop popping is enabled, label 0 (IP version 4 [IPv4] Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

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

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

```
explicit-null;
```

You can include this statement at the following hierarchy levels:

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

RELATED DOCUMENTATION

[MPLS Label Overview](#) | 419

[MPLS Label Allocation](#) | 420

Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs

By default, for both point-to-point and point-to-multipoint LSPs, penultimate-hop popping is used for MPLS traffic. MPLS labels are removed from packets on the router just before the egress router of the LSP. The plain IP packets are then forwarded to the egress router. For ultimate-hop popping, the egress router is responsible for both removing the MPLS label and processing the plain IP packet.

It can be beneficial to enable ultimate-hop popping on point-to-multipoint LSPs, particularly when transit traffic is traversing the same egress device. If you enable ultimate-hop popping, a single copy of traffic can be sent over the incoming link, saving significant bandwidth. By default, ultimate-hop popping is disabled.

You enable ultimate-hop popping for point-to-multipoint LSPs by configuring the **tunnel-services** statement. When you enable ultimate-hop popping, the Junos OS selects one of the available virtual loopback tunnel (VT) interfaces to loop back the packets to the PFE for IP forwarding. By default, the VT interface selection process is performed automatically. Bandwidth admission control is used to limit the number of LSPs that can be used on one VT interface. Once all the bandwidth is consumed on one interface, the Junos OS selects another VT interface with sufficient bandwidth for admission control.

If an LSP requires more bandwidth than is available from any of the VT interfaces, ultimate-hop popping cannot be enabled and penultimate-hop popping is enabled instead.

For ultimate-hop popping on point-to-multipoint LSPs to function properly, the egress router must have a PIC that provides tunnel services, such as the tunnel services PIC or the adaptive services PIC. Tunnel services are needed for popping the final MPLS label and for returning packets for IP address lookups.

You can explicitly configure which VT interfaces handle the RSVP traffic by including the **devices** option for the **tunnel-services** statement. The **devices** option allows you to specify which VT interfaces are to be used by RSVP. If you do not configure this option, all of the VT interfaces available to the router can be used.

To enable ultimate-hop popping for the egress point-to-multipoint LSPs on a router, configure the **tunnel-services** statement:

```
tunnel-services {
  devices device-names;
}
```

You can configure this statement at the following hierarchy levels:

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

To enable ultimate-hop popping for egress point-to-multipoint LSPs, you must also configure the **interface** statement with the **all** option:

```
interface all;
```

You must configure this statement at the [edit protocols rsvp] hierarchy level.

Tracing RSVP Protocol Traffic

To trace RSVP protocol traffic, include the **traceoptions** statement:

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


You can include this statement at the following hierarchy levels:

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

You can specify the following RSVP-specific flags in the RSVP **traceoptions** statement:

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

- **all**—All tracing operations.
- **error**—All detected error conditions
- **event**—RSVP-related events (helps to trace events related to RSVP graceful restart)
- **lmp**—RSVP-Link Management Protocol (LMP) interactions
- **packets**—All RSVP packets
- **path**—All path messages
- **pathtear**—PathTear messages
- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions, including when RSVP-signaled LSPs come up and go down.

NOTE: Use the **all** trace flag and the **detail** flag modifier with caution because these might cause the CPU to become very busy.

To view the log file generated when you enable RSVP traceoptions, issue the **show log *file-name*** command, where ***file-name*** is the file you specified using the **traceoptions** statement.

For general information about tracing and global tracing options, see the *Junos OS Routing Protocols Library*.

Examples: Tracing RSVP Protocol Traffic

Trace RSVP path messages in detail:

```
[edit]
protocols {
  rsvp {
```



```

        traceoptions {
            file rsvp size 10m files 5;
            flag path;
        }
    }
}

```

Trace all RSVP messages:

```

[edit]
protocols {
    rsvp {
        traceoptions {
            file rsvp size 10m files 5;
            flag packets;
        }
    }
}

```

Trace all RSVP error conditions:

```

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

```

Trace RSVP state transitions:

```

[edit]
protocols {
    rsvp {
        traceoptions {
            file rsvp-data;
            flag state;
        }
    }
}

```


RSVP Log File Output

The following is sample output generated by issuing the **show log file-name** command on a router on which RSVP traceoptions have been enabled with the **state** flag configured. The RSVP-signaled LSP E-D is shown being torn down on Mar 11 14:04:36.707092. On Mar 11 14:05:30.101492, it is shown coming back up.

```
user@host> show log rsvp-data
Mar 11 13:58:51 trace_on: Tracing to "/var/log/E/rsvp-data" started
Mar 11 13:58:51.286413 rsvp_iflchange for vt ifl vt-1/2/0.69206016
Mar 11 13:58:51.286718 RSVP add interface vt-1/2/0.69206016, ifindex 101, ifaddr
(null), family 1, is_appl_vt 0, already known
Mar 11 13:58:51.286818 RSVP Peer vt-1/2/0.69206016 TE-link __rpd:vt-1/2/0.69206016
Up
Mar 11 13:58:51.286978 RSVP add interface vt-1/2/0.69206016, ifindex 101, ifaddr
(null), family 3, is_appl_vt 0, already known
Mar 11 13:58:51.287962 RSVP add interface lt-1/2/0.2, ifindex 113, ifaddr (null),
family 2, is_appl_vt 0, already known
Mar 11 13:58:51.288629 RSVP add interface lt-1/2/0.2, ifindex 113, ifaddr 10.0.0.2,
family 1, is_appl_vt 0, already known
Mar 11 13:58:51.288996 RSVP add interface lt-1/2/0.17, ifindex 114, ifaddr (null),
family 2, is_appl_vt 0, already known
Mar 11 13:58:51.289593 RSVP add interface lt-1/2/0.17, ifindex 114, ifaddr (null),
family 3, is_appl_vt 0, already known
Mar 11 13:58:51.289949 RSVP add interface lt-1/2/0.17, ifindex 114, ifaddr
10.0.0.17, family 1, is_appl_vt 0, already known
Mar 11 13:58:51.290049 RSVP Peer lt-1/2/0.17 TE-link __rpd:lt-1/2/0.17 Up
Mar 11 13:59:05.042034 RSVP new bypass Bypass->10.0.0.18 on interface lt-1/2/0.17
to 10.0.0.18 avoid 0.0.0.0
Mar 11 14:04:36.707092 LSP "E-D" is Down (Reason: Reservation state deleted)
      Session: 192.168.0.4(port/tunnel ID 10321 Ext-ID 192.168.0.5) Proto 0
Sender: 192.168.0.5(port/lsp ID 1)
Mar 11 14:04:36.707661 RSVP delete resv state, session 192.168.0.4(port/tunnel ID
10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:04:36.781185 RSVP delete path state, session 192.168.0.4(port/tunnel ID
10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:04:36.781440 RSVP delete session 192.168.0.4(port/tunnel ID 10321 Ext-ID
192.168.0.5) Proto 0
Mar 11 14:05:30.101492 RSVP new Session 192.168.0.4(port/tunnel ID 10321 Ext-ID
192.168.0.5) Proto 0, session ID 3
Mar 11 14:05:30.101722 RSVP new path state, session 192.168.0.4(port/tunnel ID
10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:05:30.179124 RSVP new resv state, session 192.168.0.4(port/tunnel ID
10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:05:30.179395 RSVP PSB E-D, allocating psb resources for label 4294967295
Mar 11 14:05:30.180353 LSP "E-D" is Up
```



```
Session: 192.168.0.4(port/tunnel ID 10321 Ext-ID 192.168.0.5) Proto 0
Sender: 192.168.0.5(port/lsp ID 2)
```


Link Protection

Link protection helps to ensure that traffic going over a specific interface to a neighboring router or switch can continue to reach this router (switch) if that interface fails. When link protection is configured for an interface and an LSP that traverses this interface, a bypass LSP is created that will handle this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination. The path used can be configured explicitly, or you can rely on CSPF. The RSVP metric for the bypass LSP is set in the range of 20,000 through 29,999 (this value is not user configurable).

If a link-protected interface fails, traffic is quickly switched to the bypass LSP. Note that a bypass LSP cannot share the same egress interface with the LSPs it monitors.

In [Figure 65 on page 864](#), link protection is enabled on Interface B between Router 1 and Router 2. It is also enabled on LSP A, an LSP that traverses the link between Router 1 and Router 2. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the bypass LSP generated by link protection.

Figure 65: Link Protection Creating a Bypass LSP for the Protected Interface



Although LSPs traversing an interface can be configured to take advantage of link protection, it is important to note that it is specifically the interface that benefits from link protection. If link protection is enabled on an interface but not on a particular LSP traversing that interface, then if the interface fails, that LSP will also fail.

NOTE: Link protection does not work on unnumbered interfaces.

To protect traffic over the entire route taken by an LSP, you should configure fast reroute. For more information, see [“Configuring Fast Reroute” on page 479](#).

RELATED DOCUMENTATION

[Node Protection | 865](#)

[Multiple Bypass LSPs for Link Protection | 865](#)

[Fast Reroute, Node Protection, and Link Protection | 867](#)

[Configuring Node Protection or Link Protection for LSPs | 879](#)

Multiple Bypass LSPs for Link Protection

By default, link protection relies on a single bypass LSP to provide path protection for an interface. However, you can also specify multiple bypass LSPs to provide link protection for an interface. You can individually configure each of these bypass LSPs or create a single configuration for all of the bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.

The following algorithm describes how and when an additional bypass LSP is activated for an LSP:

1. If any currently active bypass can satisfy the requirements of the LSP (bandwidth, link protection, or node-link protection), the traffic is directed to that bypass.
2. If no active bypass LSP is available, scan through the manual bypass LSPs in first-in, first-out (FIFO) order, skipping those that are already active (each manual bypass can only be activated once). The first inactive manual bypass that can satisfy the requirements is activated and traffic is directed to that bypass.
3. If no manual bypass LSPs are available and if the **max-bypasses** statement activates multiple bypass LSPs for link protection, determine whether an automatically configured bypass LSP can satisfy the requirements. If an automatically configured bypass LSP is available and if the total number of active automatically configured bypass LSPs does not exceed the maximum bypass LSP limit (configured with the **max-bypasses** statement), activate another bypass LSP.

For information about how to configure multiple bypass LSPs for link protection, see [“Configuring Bypass LSPs” on page 872](#).

Node Protection


Node protection extends the capabilities of link protection. Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. Node protection ensures that traffic from an LSP traversing a neighboring router can continue to reach its destination even if the neighboring router fails.

When you enable node protection for an LSP, you must also enable link protection. Once enabled, node protection and link protection establish the following types of bypass LSPs:

- Next-hop bypass LSP—Provides an alternate route for an LSP to reach a neighboring router. This type of bypass LSP is established when you enable either node protection or link protection.
- Next-next-hop bypass LSP—Provides an alternate route for an LSP to get around a neighboring router en route to the destination router. This type of bypass LSP is established exclusively when node protection is configured. If a next-next-hop bypass LSP cannot be created, an attempt is made to signal a next-hop bypass LSP.

In [Figure 66 on page 866](#), node protection is enabled on Interface B on Router 1. Node protection is also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.

Figure 66: Node Protection Creating a Next-Next-Hop Bypass LSP



The time needed by node protection to switch traffic to a next-next-hop bypass LSP can be significantly longer than the time needed by link protection to switch traffic to a next-hop bypass LSP. Link protection relies on a hardware mechanism to detect a link failure, allowing it to quickly switch traffic to a next-hop bypass LSP.

Node failures are often due to software problems on the node router. Node protection relies on the receipt of hello messages from a neighboring router to determine whether it is still functioning. The time it takes node protection to divert traffic partly depends on how often the node router sends hello messages and how long it takes the node-protected router to react to having not received a hello message. However, once the failure is detected, traffic can be quickly diverted to the next-next-hop bypass LSP.

NOTE:

Node protection provides traffic protection in the event of an error or interruption of the physical link between two routers. It does not provide protection in the event of control plane errors.

The following provides an example of a control plane error:

- A transit router changes the label of a packet due to a control plane error.
- When the ingress router receives the packet, it considers the label change to be a catastrophic event and deletes both the primary LSP and the associated bypass LSP.

RELATED DOCUMENTATION

[Link Protection | 864](#)

[Fast Reroute, Node Protection, and Link Protection | 867](#)

[Configuring Node Protection or Link Protection for LSPs | 879](#)

Fast Reroute, Node Protection, and Link Protection

IN THIS SECTION

- [LSP Protection Overview | 867](#)
- [LSP Protection Types Comparison | 868](#)
- [One-to-One Backup Implementation | 868](#)
- [Facility Backup Implementation | 869](#)

This document discusses the following sections:

LSP Protection Overview

RSVP-TE extensions establish backup label-switched path (LSP) tunnels for local repair of LSP tunnels. These mechanisms enable immediate re-direction of traffic onto backup LSP tunnels, in the event of a failure.

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, describes two different types of traffic protection for RSVP-signaled LSPs:

- One-to-one backup—In this method, detour LSPs for each protected LSP is created at each potential point of local repair.
- Facility backup—In this method, a bypass tunnel is created to protect a set of LSPs that have similar backup constraints at a potential failure point, by taking advantage of the MPLS label stacking.

The one-to-one backup and the facility backup methods protect links and nodes during network failure, and can co-exist in a mixed network.

LSP Protection Types Comparison

In the Junos OS, the one-to-one backup of traffic protection is provided by fast reroute. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. This method of LSP protection cannot be shared.

In the facility backup method, the LSP traffic protection is provided on the node and link. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. Unlike fast reroute, this protecting LSP can be shared by other LSPs.

[Table 24 on page 868](#) summarizes the traffic protection types.

Table 24: One-to-One Backup Compared with Facility Backup

Comparison	One-to-One Backup	Facility Backup
Name of the protecting LSP	Detour LSP	Bypass LSP
Sharing of the protecting LSP	Cannot be shared	Can be shared by multiple LSPs
Junos configuration statements	fast-reroute	node-link-protection and link-protection

One-to-One Backup Implementation

In the one-to-one backup method, the points of local repair maintain separate backup paths for each LSP passing through a facility. The backup path terminates by merging back with the primary path at a node called the merge point. In this approach, the merge point can be any node downstream from the protected facility.

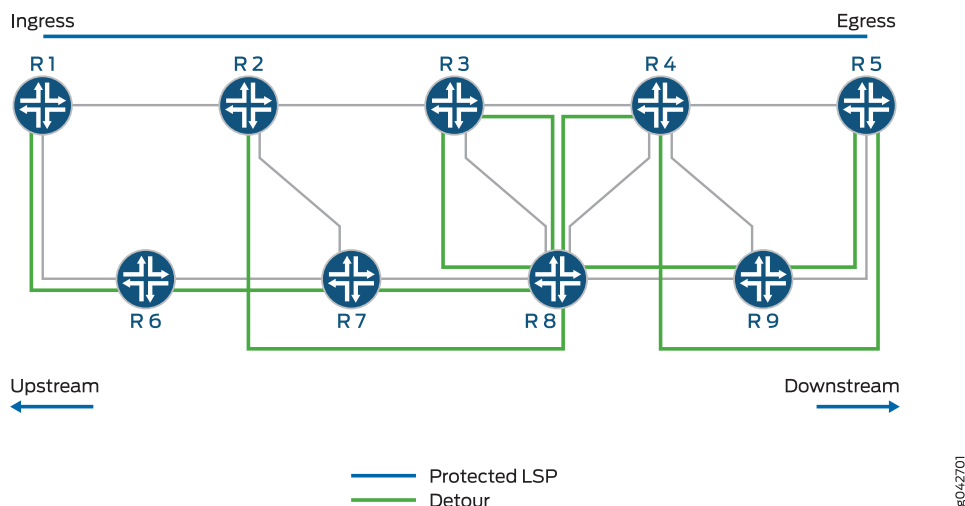
In the one-to-one backup method, an LSP is established that intersects the original LSP downstream of the point of link or node failure. A separate backup LSP is established for each LSP that is backed up.

One-to-one backup is appropriate under the following circumstances:

- Protection of a small number of LSPs relative to the total number of LSPs.
- Path selection criteria, such as bandwidth, priority, and link coloring for detour paths is critical.
- Control of individual LSPs is important.

In [Figure 67 on page 869](#), Routers R1 and R5 are the ingress and egress routers, respectively. A protected LSP is established between the two routers transiting Routers R2, R3, and R4. Router R2 provides user traffic protection by creating a partial backup LSP that merges with the protected LSP at Router R4. This partial one-to-one backup LSP is called a detour. Detours are always calculated to avoid the immediate downstream link and node, providing against both link and node failure.

Figure 67: One-to-One Backup



In the example, the protected LSP is **R1-R2-R3-R4-R5**, and the following detours are established:

- Router R1—**R1-R6-R7-R8-R3**
- Router R2—**R2-R7-R8-R4**
- Router R3—**R3-R8-R9-R5**
- Router R4—**R4-R9-R5**

To protect an LSP that traverses **N** nodes fully, there can be as many as **(N - 1)** detours. The point of local repair sends periodic refresh messages to maintain each backup path, as a result maintaining state information for backup paths protecting individual LSPs is a significant resource burden for the point of local repair. To minimize the number of LSPs in the network, it is desirable to merge a detour back to its protected LSP, when feasible. When a detour LSP intersects its protected LSP at an LSR with the same outgoing interface, it is merged.

Facility Backup Implementation

In the facility backup approach, a point of local repair maintains a single backup path to protect a set of primary LSPs traversing the point of local repair, the facility, and the merge point. The facility backup is based on interface rather than on LSP. While fast reroute protects interfaces or nodes along the entire path of a LSP, the facility backup protection can be applied on interfaces as needed. As a result, fewer states need to be maintained and refreshed which results in a scalable solution. The facility backup method is also called many-to-one backup.

The facility backup method takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created that serves to back up a set of LSPs. Such an LSP tunnel is called a bypass tunnel. In this method, a router immediately upstream from a link failure uses an alternate

interface to forward traffic to its downstream neighbor, and the merge point should be the node immediately downstream to the facility. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link. A single bypass path can safeguard a set of protected LSPs. When an outage occurs, the router immediately upstream from the link outage switches protected traffic to the bypass link, then signals the link failure to the ingress router.


The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the point of local repair. This constrains the set of LSPs being backed up through that bypass tunnel to those that pass through some common downstream nodes. All LSPs that pass through the point of local repair and through this common node, and that do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

The facility backup method is appropriate in the following situations:

- The number of LSPs to be protected is large.
- Satisfying path selection criteria (priority, bandwidth, and link coloring) for bypass paths is less critical.
- Control at the granularity of individual LSPs is not required.

In [Figure 68 on page 870](#), Routers R1 and R5 are the ingress and egress routers, respectively. Router R2 has established a bypass tunnel that protects against the failure of Router R2-R3 link and Router R3 node. A bypass tunnel is established between Routers R6 and R7. There are three different protected LSPs that are using the same bypass tunnel for protection.

Figure 68: Facility Backup



The facility backup method provides a scalability improvement, wherein the same bypass tunnel is also used to protect LSPs from any of Routers R1, R2, or R8 to any of Routers R4, R5, or R9.

RELATED DOCUMENTATION

[Fast Reroute Overview | 477](#)

[Configuring Fast Reroute | 479](#)

[Node Protection | 865](#)

[Configuring Node Protection or Link Protection for LSPs | 879](#)

Configuring Link Protection on Interfaces Used by LSPs

IN THIS SECTION

- [Configuring Bypass LSPs | 872](#)
- [Configuring Administrative Groups for Bypass LSPs | 873](#)
- [Configuring the Bandwidth for Bypass LSPs | 874](#)
- [Configuring Class of Service for Bypass LSPs | 875](#)
- [Configuring the Hop Limit for Bypass LSPs | 875](#)
- [Configuring the Maximum Number of Bypass LSPs | 876](#)
- [Disabling CSPF for Bypass LSPs | 876](#)
- [Disabling Node Protection for Bypass LSPs | 877](#)
- [Configuring the Optimization Interval for Bypass LSPs | 877](#)
- [Configuring an Explicit Path for Bypass LSPs | 878](#)
- [Configuring the Amount of Bandwidth Subscribed for Bypass LSPs | 878](#)
- [Configuring Priority and Preemption for Bypass LSPs | 879](#)

When you configure node protection or link protection on a router for LSPs as described in “[Configuring Node Protection or Link Protection for LSPs](#)” on page 879, you also must configure the **link-protection** statement on the RSVP interfaces used by the LSPs.

To configure link protection on the interfaces used by the LSPs, include the **link-protection** statement:

```
link-protection {
  disable;
  admin-group
    exclude group-names;
    include-all group-names;
```



```

    include-any group-names;
}
bandwidth bps;
bypass bypass-name {
    bandwidth bps;
    description text;
    hop-limit number;
    no-cspf;
    path address <strict | loose>;
    priority setup-priority reservation-priority;
    to address;
}
class-of-service cos-value;
hop-limit number;
max-bypasses number;
no-cspf;
no-node-protection;
optimize-timer seconds;
path address <strict | loose>;
priority setup-priority reservation-priority;
subscription percent {
    ct0 percent;
    ct1 percent;
    ct2 percent;
    ct3 percent;
}
}

```

You can include this statement at the following hierarchy levels:

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

All the statements under **link-protection** are optional.

The following sections describe how to configure link protection:

Configuring Bypass LSPs

You can configure specific bandwidth and path constraints for a bypass LSP. Each manual bypass LSP on a router should have a unique “to” IP address. You can also individually configure each bypass LSP generated when you enable multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints (if any).

If you specify the **bandwidth**, **hop-limit**, and **path** statements for the bypass LSP, these values take precedence over the values configured at the [edit protocols rsvp interface *interface-name* link-protection] hierarchy level. The other attributes (**subscription**, **no-node-protection**, and **optimize-timer**) are inherited from the general constraints.

To configure a bypass LSP, specify a name for the bypass LSP using the **bypass** statement. The name can be up to 64 characters in length.

```
bypass bypass-name {
  bandwidth bps;
  description text;
  class-of-service cos-value;
  hop-limit number;
  no-cspf;
  path address <strict | loose>;
  priority setup-priority reservation-priority;
  to address;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass LSPs

If you configure a bypass LSP, you must also configure the **to** statement. The **to** statement specifies the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.

Configuring Administrative Groups for Bypass LSPs

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups. You can configure administrative groups for bypass LSPs. For more information about configuring administrative groups, see [“Configuring Administrative Groups for LSPs” on page 532](#).

To configure administrative groups for bypass LSPs, include the **admin-group** statement:

```
admin-group {
  exclude group-names;
```



```

include-all group-names;
include-any group-names;
}

```

To configure an administrative group for all of the bypass LSPs, include the **admin-group** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

To configure an administrative groups for a specific bypass LSP, include the **admin-group** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Bandwidth for Bypass LSPs

You can specify the amount of bandwidth allocated for automatically generated bypass LSPs or you can individually specify the amount of bandwidth allocated for each LSP.

If you have enabled multiple bypass LSPs, this statement is required.

To specify the bandwidth allocation, include the **bandwidth** statement:

```

bandwidth bps;

```

For automatically generated bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring Class of Service for Bypass LSPs

You can specify the class-of-service value for bypass LSPs by including the **class-of-service** statement:

```
class-of-service cos-value;
```

To apply a class-of-service value to all the automatically generated bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

To configure a class-of-service value for a specific bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Hop Limit for Bypass LSPs

You can specify the maximum number of hops a bypass can traverse. By default, each bypass can traverse a maximum of 255 hops (the ingress and egress routers count as one hop each, so the minimum hop limit is two).

To configure the hop limit for bypass LSPs, include the **hop-limit** statement:

```
hop-limit number;
```

For automatically generated bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Maximum Number of Bypass LSPs

You can specify the maximum number of dynamic bypass LSPs permitted for protecting an interface using the **max-bypasses** statement at the **[edit protocols rsvp interface *interface-name* link-protection]** hierarchy level. When this statement is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled.

By default, this option is disabled and only one bypass is enabled for each interface. You can configure a value of between **0** through **99** for the **max-bypasses** statement. Configuring a value of **0** prevents the creation of any dynamic bypass LSPs for the interface. If you configure a value of **0** for the **max-bypasses** statement, you need to configure one or more static bypass LSPs to enable link protection on the interface.

If you configure the **max-bypasses** statement, you must also configure the **bandwidth** statement (discussed in [“Configuring the Bandwidth for Bypass LSPs” on page 874](#)).

To configure the maximum number of bypass LSPs for a protected interface, include the **max-bypasses** statement:

```
max-bypasses number;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp *interface interface-name* link-protection]**
- **[edit logical-systems *logical-system-name* protocols rsvp *interface interface-name* link-protection]**

Disabling CSPF for Bypass LSPs

Under certain circumstances, you might need to disable CSPF computation for bypass LSPs and use the configured Explicit Route Object (ERO) if available. For example, a bypass LSP might need to traverse multiple OSPF areas or IS-IS levels, preventing the CSPF computation from working. To ensure that link and node protection function properly in this case, you have to disable CSPF computation for the bypass LSP.

You can disable CSPF computation for all bypass LSPs or for specific bypass LSPs.

To disable CSPF computation for bypass LSPs, include the **no-cspf** statement:

```
no-cspf;
```

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

Disabling Node Protection for Bypass LSPs

You can disable node protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

To disable node protection for bypass LSPs, include the **no-node-protection** statement:

```
no-node-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

Configuring the Optimization Interval for Bypass LSPs

You can configure an optimization interval for bypass LSPs using the **optimize-timer** statement. At the end of this interval, an optimization process is initiated that attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all of the bypasses, or both. You can configure an optimization interval from 1 through 65,535 seconds. A default value of 0 disables bypass LSP optimization.

When you configure the **optimize-timer** statement, bypass LSPs are reoptimized automatically when you configure or change the configuration of any of the following:

- Administrative group for a bypass LSP—The configuration for an administrative group has been changed on a link along the path used by the bypass LSP. Configure an administrative group using the **admin-group** statement at the [edit protocols rsvp **interface** *interface-name* link-protection] hierarchy level.
- Fate sharing group—The configuration for a fate sharing group has been changed. Configure a fate sharing group using the **group** statement at the [edit routing-options fate-sharing] hierarchy level.
- IS-IS overload—The configuration for IS-IS overload has been changed on a router along the path used by the bypass LSP. Configure IS-IS overload using the **overload** statement at the [edit protocols isis] hierarchy level.
- IGP metric—The IGP metric has been changed on a link along the path used by the bypass LSP.

To configure the optimization interval for bypass LSPs, include the **optimize-timer** statement:

```
optimize-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]

- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

Configuring an Explicit Path for Bypass LSPs

By default, when you establish a bypass LSP to an adjacent neighbor, CSPF is used to discover the least-cost path. The **path** statement allows you to configure an explicit path (a sequence of strict or loose routes), giving you control over where and how the bypass LSP is established. To configure an explicit path, include the **path** statement:

```
path address <strict | loose>;
```

For automatically generated bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

Configuring the Amount of Bandwidth Subscribed for Bypass LSPs

You can configure the amount of bandwidth subscribed to bypass LSPs. You can configure the bandwidth subscription for the whole bypass LSP or for each class type that might traverse the bypass LSP. You can configure any value between 1 percent and 65,535 percent. By configuring a value less than 100 percent, you are undersubscribing the bypass LSPs. By configuring a value greater than 100 percent, you are oversubscribing the bypass LSPs.

The ability to oversubscribe the bandwidth for the bypass LSPs makes it possible to more efficiently use network resources. You can configure the bandwidth for the bypass LSPs based on the average network load as opposed to the peak load.

To configure the amount of bandwidth subscribed for bypass LSPs, include the **subscription** statement:

```
subscription percentage {
  ct0 percentage;
  ct1 percentage;
  ct2 percentage;
  ct3 percentage;
}
```


You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

Configuring Priority and Preemption for Bypass LSPs

When there is insufficient bandwidth to establish a more important LSP, you might want to tear down a less important existing LSP to release the bandwidth. You do this by preempting the existing LSP.

For more detailed information on configuring setup priority and reservation priority for LSPs, see [“Configuring Priority and Preemption for LSPs” on page 531](#).

To configure the bypass LSP’s priority and preemption properties, include the **priority** statement:

```
priority setup-priority reservation-priority;
```

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

Configuring Node Protection or Link Protection for LSPs

When you configure node protection or link protection on a router or switch, bypass LSPs are created to the next-hop or next-next-hop routers (switches) for the LSPs traversing the router (switch). You must configure node protection or link protection for each LSP that you want protected. To extend protection along the entire path used by an LSP, you must configure protection on each router that the LSP traverses.

You can configure node protection or link protection for both static and dynamic LSPs.

To configure node protection on a router for a specified LSP, include the **node-link-protection** statement:

```
node-link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

To configure link protection on a router for a specified LSP, include the **link-protection** statement:

```
link-protection;
```


You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]

NOTE: To complete the configuration of node or link protection, you must also configure link protection on all unidirectional RSVP interfaces that the LSPs traverse, as described in [“Configuring Link Protection on Interfaces Used by LSPs”](#) on page 871.

RELATED DOCUMENTATION

[Configuring Link Protection on Interfaces Used by LSPs | 871](#)

[Link Protection | 864](#)

[Node Protection | 865](#)

[Multiple Bypass LSPs for Link Protection | 865](#)

[Fast Reroute, Node Protection, and Link Protection | 867](#)

Configuring Inter-AS Node and Link Protection

To interoperate with other vendors' equipment, the Junos OS supports the record route object (RRO) node ID subobject for use in inter-AS link and node protection configurations. The RRO node ID subobject is defined in RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*. This functionality is enabled by default in Junos OS Release 9.4 and later.

If you have Juniper Networks routers running Junos OS Release 9.4 and later releases in the same MPLS-TE network as routers running Junos OS Release 8.4 and earlier releases, you might need to disable the RRO node ID subobject by configuring the **no-node-id-subobject** statement:

```
no-node-id-subobject;
```

You can include this statement at the following hierarchy levels:

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

RSVP Graceful Restart

RSVP graceful restart allows a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

RSVP graceful restart is described in the following sections:

- [RSVP Graceful Restart Standard on page 881](#)
- [RSVP Graceful Restart Terminology on page 881](#)
- [RSVP Graceful Restart Operation on page 882](#)
- [Processing the Restart Cap Object on page 883](#)

RSVP Graceful Restart Standard

RSVP graceful restart is described in RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (only Section 9, “Fault Handling”).

RSVP Graceful Restart Terminology

R

Recovery time (in milliseconds)	Applies only when the control channel is up (the hello exchange is complete) before the restart time. Applies only to nodal faults.
	When a graceful restart is in progress, the time left to complete a recovery is advertised. At other times, this value is zero. The maximum advertised recovery time is 2 minutes (120,000 milliseconds).
	During the recovery time, a restarting node attempts to recover its lost states with assistance from its neighbors. The neighbor of the restarting node must send the path messages with the recovery labels to the restarting node within a period of one-half the recovery time. The restarting node considers its graceful restart complete after its advertised recovery time.

Restart time	The default value is 60,000 milliseconds (1 minute). The restart time is advertised in the hello message.
(in milliseconds)	The time indicates how long a neighbor should wait to receive a hello message from a restarting router before declaring that router dead and purging states.

The Junos OS can override a neighbor's advertised restart time if the time is greater than one-third the local restart time. For example, given the default restart time of 60 seconds, a router would wait 20 seconds or less to receive a hello message from a restarting neighbor. If the restart time is zero, the restarting neighbor can immediately be declared dead.

RSVP Graceful Restart Operation

For RSVP graceful restart to function, the feature must be enabled on the global routing instance. RSVP graceful restart can be disabled at the protocol level (for RSVP alone) or at the global level for all protocols.

RSVP graceful restart requires the following of a restarting router and the router's neighbors:

- For the restarting router, RSVP graceful restart attempts to maintain the routes installed by RSVP and the allocated labels, so that traffic continues to be forwarded without disruption. RSVP graceful restart is done quickly enough to reduce or eliminate the impact on neighboring nodes.
- The neighboring routers must have RSVP graceful restart helper mode enabled, thus allowing them to assist a router attempting to restart RSVP.

An object called Restart Cap that is sent in RSVP hello messages advertises a node's restart capability. The neighboring node sends a Recover Label object to the restarting node to recover its forwarding state. This object is essentially the old label that the restarting node advertised before the node went down.

The following lists the RSVP graceful restart behaviors, which vary depending on the configuration and on which features are enabled:

- If you disable helper mode, the Junos OS does not attempt to help a neighbor restart RSVP. Any information that arrives with a Restart Cap object from a neighbor is ignored.
- When you enable graceful restart under the routing instance configuration, the router can restart gracefully with the help of its neighbors. RSVP advertises a Restart Cap object (RSVP RESTART) in hello messages in which restart and recovery times are specified (neither value is 0).
- If you explicitly disable RSVP graceful restart under the **[protocols rsvp]** hierarchy level, the Restart Cap object is advertised with restart and recovery times specified as 0. The restart of neighboring routers is supported (unless helper mode is disabled), but the router itself does not preserve the RSVP forwarding state and cannot recover its control state.

- If after a restart RSVP realizes that no forwarding state has been preserved, the Restart Cap object is advertised with restart and recovery times specified as 0.
- If graceful restart and helper mode are disabled, RSVP graceful restart is completely disabled. The router neither recognizes nor advertises the RSVP graceful restart objects.

You cannot explicitly configure values for the restart and recovery times.

Unlike other protocols, there is no way for RSVP to determine that it has completed a restart procedure, other than a fixed timeout. All RSVP graceful restart procedures are timer-based. A **show rsvp version** command might indicate that the restart is still in progress even if all RSVP sessions are back up and the routes are restored.

Processing the Restart Cap Object

The following assumptions are made about a neighbor based on the Restart Cap object (assuming that a control channel failure can be distinguished unambiguously from a node restart):

- A neighbor that does not advertise the Restart Cap object in its hello messages cannot assist a router with state or label recovery, nor can it perform an RSVP graceful restart.
- After a restart, a neighbor advertising a Restart Cap object with a restart time equal to any value and a recovery time equal to 0 has not preserved its forwarding state. When a recovery time equals 0, the neighbor is considered dead and any states related to this neighbor are purged, regardless of the value of the restart time.
- After a restart, a neighbor advertising its recovery time with a value other than 0 can keep or has kept the forwarding state. If the local router is helping its neighbor with restart or recovery procedures, it sends a Recover Label object to this neighbor.

Configuring RSVP Graceful Restart

IN THIS SECTION

- [Enabling Graceful Restart for All Routing Protocols | 884](#)
- [Disabling Graceful Restart for RSVP | 884](#)
- [Disabling RSVP Helper Mode | 885](#)
- [Configuring the Maximum Helper Recovery Time | 885](#)
- [Configuring the Maximum Helper Restart Time | 885](#)

The following RSVP graceful restart configurations are possible:

- Graceful restart and helper mode are both enabled (the default).
- Graceful restart is enabled but helper mode is disabled. A router configured in this way can restart gracefully, but cannot help a neighbor with its restart and recovery procedures.
- Graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully, but can help a restarting neighbor.
- Graceful restart and helper mode both are disabled. This configuration completely disables RSVP graceful restart (including restart and recovery procedures and helper mode). The router behaves like a router that does not support RSVP graceful restart.

NOTE: In order to turn on RSVP graceful restart, you must set the global graceful restart timer to at least 180 seconds.

The following sections describe how to configure RSVP graceful restart:

Enabling Graceful Restart for All Routing Protocols

To enable graceful restart for RSVP, you need to enable graceful restart for all the protocols that support graceful restart on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library*.

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

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

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

Disabling Graceful Restart for RSVP

By default, RSVP graceful restart and RSVP helper mode are enabled when you enable graceful restart. However, you can disable one or both of these capabilities.

To disable RSVP graceful restart and recovery, include the **disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level:

```
disable;
```

Disabling RSVP Helper Mode

To disable RSVP helper mode, include the **helper-disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level:

```
helper-disable;
```

Configuring the Maximum Helper Recovery Time

To configure the amount of time the router retains the state of its RSVP neighbors while they undergo a graceful restart, include the **maximum-helper-recovery-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to recover.

```
maximum-helper-recovery-time seconds;
```

Configuring the Maximum Helper Restart Time

To configure the delay between when the router discovers that a neighboring router has gone down and when it declares the neighbor down, include the **maximum-helper-restart-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to restart.

```
maximum-helper-restart-time seconds;
```

RSVP LSP Tunnels Overview

A Resource Reservation Protocol (RSVP) label-switched path (LSP) tunnel enables you to send RSVP LSPs inside other RSVP LSPs. This enables a network administrator to provide traffic engineering from one end of the network to the other. A useful application for this feature is to connect customer edge (CE) routers with provider edge (PE) routers by using an RSVP LSP, and then tunnel this edge LSP inside a second RSVP LSP traveling across the network core.

You should have a general understanding of MPLS and label switching concepts. For more information about MPLS, see the *Junos MPLS Applications Configuration Guide*.

An RSVP LSP tunnel adds the concept of a forwarding adjacency, similar to the one used for generalized Multiprotocol Label Switching (GMPLS). (For more information about GMPLS, see *Junos GMPLS User Guide*.)

The forwarding adjacency creates a tunneled path for sending data between peer devices in an RSVP LSP network. Once a forwarding adjacency LSP (FA-LSP) has been established, other LSPs can be sent over the FA-LSP by using Constrained Shortest Path First (CSPF), Link Management Protocol (LMP), Open Shortest Path First (OSPF), and RSVP.

To enable an RSVP LSP tunnel, the Junos OS uses the following mechanisms:

- LMP—Originally designed for GMPLS, LMP establishes forwarding adjacencies between RSVP LSP tunnel peers, and maintains and allocates resources for traffic engineering links.
- OSPF extensions—OSPF was designed to route packets to physical and logical interfaces related to a Physical Interface Card (PIC). This protocol has been extended to route packets to virtual peer interfaces defined in an LMP configuration.
- RSVP-TE extensions—RSVP-TE was designed to signal the setup of packet LSPs to physical interfaces. The protocol has been extended to request path setup for packet LSPs traveling to virtual peer interfaces defined in an LMP configuration.

NOTE: Beginning with Junos OS Release 15.1, multi-instance support is extended to MPLS RSVP-TE. This support is available only for virtual router instance type. A router can create and participate in multiple independent TE topology partitions, which allows each partitioned TE domain to scale independently. Multi-instance RSVP-TE provides the flexibility to hand pick the control-plane entities that need to be instance-aware, for example, a router can participate in multiple TE instances while still running a single BGP instance.

The Junos OS implementation of MPLS RSVP-TE is scaled to enhance the usability, visibility, configuration, and troubleshooting of LSPs in Junos OS Release 16.1.

These enhancements make the RSVP-TE configuration easier at scale by:

- Ensuring the LSP data-plane readiness during LSP resignaling before traffic traverses the LSP with the RSVP-TE LSP self-ping mechanism.

An LSP should not start to carry traffic unless it is known to have been programmed in the data plane. Data exchange in the LSP data plane, such as LSP ping requests, happens at the ingress router before switching traffic to an LSP, or to its MBB instance. In large networks, this traffic can overwhelm an LSP egress router, as the egress LSP needs to respond to the LSP ping requests. The LSP self-ping mechanism enables the ingress LER to create LSP ping response messages and send them over the LSP data plane. On receiving these messages, the egress LER forwards them to the ingress, indicating the liveness of the LSP data plane. This ensures that the LSP does not start to carry traffic before the data plane has been programmed.

- Removing the current hard limit of 64K LSPs on an ingress router and scaling the total number of LSPs with RSVP-TE signaled LSPs. There can be up to 64K LSPs configured on a per-egress basis. Earlier, this limit was the aggregate number of LSPs that could be configured on the ingress LER.
- Preventing abrupt tearing down of LSPs by the ingress router because of delay in signaling the LSP at the transit routers.
- Enabling a flexible view of LSP data-sets to facilitate LSP characteristic data visualization.

NOTE: Starting with Junos OS Release 17.4, a default timer of 1800 seconds for self-ping is introduced.

The following limitations exist for LSP hierarchies:

- Circuit cross-connect (CCC)-based LSPs are not supported.
- Graceful restart is not supported.
- Link protection is not available for FA-LSPs or at the egress point of the forwarding adjacency.
- Point-to-multipoint LSPs are not supported across FA-LSPs.

RELATED DOCUMENTATION

[Example: RSVP LSP Tunnel Configuration](#) | 888

Example: RSVP LSP Tunnel Configuration

Figure 69: RSVP LSP Tunnel Topology Diagram




Figure 69 on page 888 shows an end-to-end RSVP LSP called **e2e_lsp_r0r5** that originates on Router 0 and terminates on Router 5. In transit, this LSP traverses the FA-LSP **fa_lsp_r1r4**. The return path is represented by the end-to-end RSVP LSP **e2e_lsp_r5r0** that travels over the FA-LSP **fa_lsp_r4r1**.

On Router 0, configure the end-to-end RSVP LSP that travels to Router 5. Use a strict path that traverses Router 1 and the LMP traffic engineering link traveling from Router 1 to Router 4.

Router 0

```
[edit]
interfaces {
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.2.1/30;
      }
      family mpls;
    }
  }
}
```



```

lo0 {
  unit 0 {
    family inet {
      address 10.255.41.222/32;
    }
    family mpls;
  }
}
routing-options {
  forwarding-table {
    export pplib;
  }
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    admin-groups {
      fa 1;
      backup 2;
      other 3;
    }
    label-switched-path e2e_lsp_r0r5 { # An end-to-end LSP traveling to Router 5.
      to 10.255.41.221;
      bandwidth 30k;
      primary path-fa; # Reference the requested path here.
    }
    path path-fa { # Configure the strict path here.
      10.1.2.2 strict;
      172.16.30.2 strict; # This traverses the TE link heading to Router 4.
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface so-3/2/1.0 {
      admin-group other;

```



```

    }
    interface so-0/0/3.0 {
        admin-group other;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
    }
}
}
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
}

```

On Router 1, configure an FA-LSP to reach Router 4. Establish an LMP traffic engineering link and LMP peer relationship with Router 4. Reference the FA-LSP in the traffic engineering link and add the peer interface into both OSPF and RSVP.

When the return path end-to-end LSP arrives at Router 1, the routing platform performs a routing lookup and can forward traffic to Router 0. Make sure you configure OSPF correctly between Routers 0 and 1.

Router 1

```

[edit]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.2.3.1/30;
            }
            family mpls;
        }
    }
}

```



```

    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.2.4.1/30;
        }
        family mpls;
    }
}
so-0/0/3 {
    unit 0 {
        family inet {
            address 10.1.2.2/30;
        }
        family mpls;
    }
}
fe-0/1/2 {
    unit 0 {
        family inet {
            address 10.2.5.1/30;
        }
        family mpls;
    }
}
at-1/0/0 {
    atm-options {
        vpi 1;
    }
    unit 0 {
        vci 1.100;
        family inet {
            address 10.2.3.5/30;
        }
        family mpls;
    }
}
}
routing-options {
    forwarding-table {
        export [ pplb choose_lsp ];
    }
}

```



```

    }
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
        peer-interface r4; # Apply the LMP peer interface here.
    }
    mpls {
        admin-groups {
            fa 1;
            backup 2;
            other 3;
        }
        label-switched-path fa_lsp_r1r4 { # Configure your FA-LSP to Router 4 here.
            to 10.255.41.217;
            bandwidth 400k;
            primary path_r1r4; # Apply the FA-LSP path here.
        }
        path path_r1r4 { # Configure the FA-LSP path here.
            10.2.4.2;
            10.4.5.2;
            10.3.5.1;
        }
        interface so-0/0/3.0 {
            admin-group other;
        }
        interface so-0/0/1.0 {
            admin-group fa;
        }
        interface at-1/0/0.0 {
            admin-group backup;
        }
        interface fe-0/1/2.0 {
            admin-group backup;
        }
        interface so-0/0/2.0 {
            admin-group fa;
        }
    }
}

```



```

ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
        peer-interface r4; # Apply the LMP peer interface here.
    }
}

link-management { # Configure LMP statements here.
    te-link link_r1r4 { # Assign a name to the TE link here.
        local-address 172.16.30.1; # Configure a local address for the TE link.
        remote-address 172.16.30.2; # Configure a remote address for the TE link.
        te-metric 1; # Manually set a metric here if you are not relying on CSPF.
        label-switched-path fa_lsp_r1r4; # Reference the FA-LSP here.
    }
    peer r4 { # Configure LMP peers here.
        address 10.255.41.217; # Configure the loopback address of your peer here.
        te-link link_r1r4; # Apply the LMP TE link here.
    }
}

policy-options {
    policy-statement choose_lsp {
        term A {
            from community choose_e2e_lsp;
            then {
                install-nexthop strict lsp e2e_lsp_r1r4;
                accept;
            }
        }
        term B {
            from community choose_fa_lsp;
            then {
                install-nexthop strict lsp fa_lsp_r1r4;
                accept;
            }
        }
    }
}

policy-statement pplb {
    then {

```



```

        load-balance per-packet;
    }
}
community choose_e2e_lsp members 1000:1000;
community choose_fa_lsp members 2000:2000;
community set_e2e_lsp members 1000:1000;
community set_fa_lsp members 2000:2000;
}

```

On Router 2, configure OSPF, MPLS, and RSVP on all interfaces that transport the FA-LSPs across the core network.

Router 2

```

[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.4.5.1/30;
      }
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.4.2/30;
      }
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.2.4.2/30;
      }
      family mpls;
    }
  }
}

```



```

fe-0/1/2 {
  unit 0 {
    family inet {
      address 10.3.4.2/30;
    }
    family mpls;
  }
}
}
routing-options {
  forwarding-table {
    export pplb;
  }
}
protocols { # OSPF, MPLS, and RSVP form the core backbone for the FA-LSPs.
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    admin-groups {
      fa 1;
      backup 2;
      other 3;
    }
    path path_r1 {
      10.2.4.1;
    }
    path path_r3r4 {
      10.4.5.2;
      10.3.5.1;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
    interface so-0/0/1.0 {
      admin-group other;
    }
    interface fe-0/1/2.0 {

```



```

        admin-group backup;
    }
    interface so-0/0/2.0 {
        admin-group fa;
    }
    interface so-0/0/0.0 {
        admin-group fa;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
    }
}
}
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
}

```

On Router 3, configure OSPF, MPLS, and RSVP on all interfaces that transport the FA-LSPs across the core network.

Router 3

```

[edit]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.4.5.2/30;
            }
            family mpls;
        }
    }
}

```



```

    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.5.6.1/30;
        }
        family mpls;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.3.5.2/30;
        }
        family mpls;
    }
}
fe-0/1/2 {
    unit 0 {
        family inet {
            address 10.2.5.2/30;
        }
        family mpls;
    }
}
}
routing-options {
    forwarding-table {
        export pplb;
    }
}
}
protocols { # OSPF, MPLS, and RSVP form the core backbone for the FA-LSPs.
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        admin-groups {
            fa 1;

```



```

        backup 2;
        other 3;
    }
    path path_r4 {
        10.3.5.1;
    }
    path path_r2r1 {
        10.4.5.1;
        10.2.4.1;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface so-0/0/2.0 {
        admin-group fa;
    }
    interface fe-0/1/2.0 {
        admin-group backup;
    }
    interface so-0/0/1.0 {
        admin-group other;
    }
    interface so-0/0/0.0 {
        admin-group fa;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
    }
}
}
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}

```



```

    }
}

```

On Router 4, configure a return path FA-LSP to reach Router 1. Establish an LMP traffic engineering link and LMP peer relationship with Router 1. Reference the FA-LSP in the traffic engineering link and add the peer interface into both OSPF and RSVP.

When the initial end-to-end LSP arrives at Router 4, the routing platform performs a routing lookup and can forward traffic to Router 5. Make sure you configure OSPF correctly between Router 4 and Router 5.

Router 4

```

[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.3.6.1/30;
      }
      family mpls;
    }
  }
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.2.3.2/30;
      }
      family mpls;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.3.5.1/30;
      }
      family mpls;
    }
  }
  fe-0/1/2 {
    unit 0 {

```



```

        family inet {
            address 10.3.4.1/30;
        }
        family mpls;
    }
}
at-1/0/0 {
    atm-options {
        vpi 1;
    }
    unit 0 {
        vci 1.100;
        family inet {
            address 10.2.3.6/30;
        }
        family mpls;
    }
}
}
routing-options {
    forwarding-table {
        export [ pplb choose_lsp ];
    }
}
protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
        peer-interface r1; # Apply the LMP peer interface here.
    }
    mpls {
        admin-groups {
            fa 1;
            backup 2;
            other 3;
        }
        label-switched-path fa_lsp_r4r1 { # Configure your FA-LSP here.
            to 10.255.41.216;
            bandwidth 400k;
            primary path_r4r1; # Apply the FA-LSP path here.
        }
    }
}

```



```

}
path path_r4r1 { # Configure the FA-LSP path here.
    10.3.5.2;
    10.4.5.1;
    10.2.4.1;
}
interface all;
interface fxp0.0 {
    disable;
}
interface at-1/0/0.0 {
    admin-group backup;
}
interface so-0/0/2.0 {
    admin-group fa;
}
interface fe-0/1/2.0 {
    admin-group backup;
}
interface so-0/0/0.0 {
    admin-group other;
}
interface so-0/0/1.0 {
    admin-group fa;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
        peer-interface r1; # Apply the LMP peer interface here.
    }
}
link-management { # Configure LMP statements here.
    te-link link_r4r1 { # Assign a name to the TE link here.
        local-address 172.16.30.2; # Configure a local address for the TE link.
        remote-address 172.16.30.1; # Configure a remote address for the TE link.
        te-metric 1; # Manually set a metric here if you are not relying on CSPF.
        label-switched-path fa_lsp_r4r1; # Reference the FA-LSP here.
    }
}

```



```

    }
    peer r1 { # Configure LMP peers here.
        address 10.255.41.216; # Configure the loopback address of your peer here.
        te-link link_r4r1; # Apply the LMP TE link here.
    }
}
policy-options {
    policy-statement choose_lsp {
        term A {
            from community choose_e2e_lsp;
            then {
                install-nexthop strict lsp e2e_lsp_r4r1;
                accept;
            }
        }
        term B {
            from community choose_fa_lsp;
            then {
                install-nexthop strict lsp fa_lsp_r4r1;
                accept;
            }
        }
    }
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
    community choose_e2e_lsp members 1000:1000;
    community choose_fa_lsp members 2000:2000;
    community set_e2e_lsp members 1000:1000;
    community set_fa_lsp members 2000:2000;
}

```

On Router 5, configure the return path end-to-end RSVP LSP that travels to Router 0. Use a strict path that traverses Router 4 and the LMP traffic engineering link traveling from Router 4 to Router 1.

Router 5


```

[edit]
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.3.6.2/30;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.41.221/32;
      }
    }
  }
}
routing-options {
  forwarding-table {
    export pplb;
  }
}
protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    admin-groups {
      fa 1;
      backup 2;
      other 3;
    }
    label-switched-path e2e_lsp_r5r0 { # An end-to-end LSP returning to Router 0.
      to 10.255.41.222;
      bandwidth 30k;
      primary path-fa; # Reference the requested path here.
    }
    path path-fa { # Configure the strict path here.

```



```

    10.3.6.1 strict;
    172.16.30.1 strict; # This traverses the TE link heading to Router 1.
}
interface all;
interface fxp0.0 {
    disable;
}
interface so-0/0/2.0 {
    admin-group other;
}
interface so-0/0/1.0 {
    admin-group other;
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fxp0.0 {
            disable;
        }
        interface all;
    }
}
}
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
}

```

Verifying Your Work

IN THIS SECTION

- Router 0 | 905
- Router 1 | 911

To verify that your RSVP LSP tunnel is working correctly, issue the following commands:

- **show ted database (extensive)**
- **show rsvp session name (extensive)**
- **show link-management**
- **show link-management te-link name (detail)**

To see these commands used with the configuration example, see the following sections:

Router 0

On Router 0, you can verify that the FA-LSPs appear as valid paths in the traffic engineering database. In this case, look for the paths from Router 1 (**10.255.41.216**) and Router 4 (**10.255.41.217**) that reference the LMP traffic engineering link addresses of **172.16.30.1** and **172.16.30.2**. You can also issue the **show rsvp session extensive** command to look for the path of the end-to-end LSP as it travels to Router 5 over the FA-LSP.

```
user@router0> show ted database
```

```
TED database: 0 ISIS nodes 8 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.214                    Rtr    486     4     4 OSPF(0.0.0.0)
  To: 10.255.41.222, Local: 10.1.4.2, Remote: 10.1.4.1
  To: 10.255.41.216, Local: 10.2.4.2, Remote: 10.2.4.1
  To: 10.255.41.215, Local: 10.4.5.1, Remote: 10.4.5.2
  To: 10.3.4.1-1, Local: 10.3.4.2, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.215                    Rtr    187     4     4 OSPF(0.0.0.0)
  To: 10.255.41.214, Local: 10.4.5.2, Remote: 10.4.5.1
  To: 10.255.41.217, Local: 10.3.5.2, Remote: 10.3.5.1
  To: 10.255.41.221, Local: 10.5.6.1, Remote: 10.5.6.2
  To: 10.2.5.1-1, Local: 10.2.5.2, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.216                   Rtr    396     6     6 OSPF(0.0.0.0)
  To: 10.255.41.222, Local: 10.1.2.2, Remote: 10.1.2.1
  To: 10.255.41.214, Local: 10.2.4.1, Remote: 10.2.4.2
  To: 10.255.41.217, Local: 10.2.3.1, Remote: 10.2.3.2
  To: 10.255.41.217, Local: 172.16.30.1, Remote: 172.16.30.2
  To: 10.255.41.217, Local: 10.2.3.5, Remote: 10.2.3.6
  To: 10.2.5.1-1, Local: 10.2.5.1, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.217                   Rtr    404     6     6 OSPF(0.0.0.0)
  To: 10.255.41.216, Local: 10.2.3.2, Remote: 10.2.3.1
  To: 10.255.41.216, Local: 172.16.30.2, Remote: 172.16.30.1
```



```

To: 10.255.41.216, Local: 10.2.3.6, Remote: 10.2.3.5
To: 10.255.41.215, Local: 10.3.5.1, Remote: 10.3.5.2
To: 10.255.41.221, Local: 10.3.6.1, Remote: 10.3.6.2
To: 10.3.4.1-1, Local: 10.3.4.1, Remote: 0.0.0.0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.221                    Rtr   481     2     2 OSPF(0.0.0.0)
    To: 10.255.41.215, Local: 10.5.6.2, Remote: 10.5.6.1
    To: 10.255.41.217, Local: 10.3.6.2, Remote: 10.3.6.1
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.41.222                    Rtr  2883     2     2 OSPF(0.0.0.0)
    To: 10.255.41.216, Local: 10.1.2.1, Remote: 10.1.2.2
    To: 10.255.41.214, Local: 10.1.4.1, Remote: 10.1.4.2

user@router0> show ted database 10.255.41.216 extensive
TED database: 0 ISIS nodes 8 INET nodes
NodeID: 10.255.41.216
  Type: Rtr, Age: 421 secs, LinkIn: 6, LinkOut: 6
  Protocol: OSPF(0.0.0.0)
    To: 10.255.41.222, Local: 10.1.2.2, Remote: 10.1.2.1
      Color: 0x8 other
      Metric: 1
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.4Mbps   [1] 155.4Mbps   [2] 155.4Mbps   [3] 155.4Mbps
        [4] 155.4Mbps   [5] 155.4Mbps   [6] 155.4Mbps   [7] 155.4Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 155.4Mbps   [1] 155.4Mbps   [2] 155.4Mbps   [3] 155.4Mbps
          [4] 155.4Mbps   [5] 155.4Mbps   [6] 155.4Mbps   [7] 155.4Mbps
    To: 10.255.41.214, Local: 10.2.4.1, Remote: 10.2.4.2
      Color: 0x2 fa
      Metric: 1
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.12Mbps   [1] 155.12Mbps   [2] 155.12Mbps   [3] 155.12Mbps
        [4] 155.12Mbps   [5] 155.12Mbps   [6] 155.12Mbps   [7] 155.12Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:

```



```

        [0] 155.12Mbps    [1] 155.12Mbps    [2] 155.12Mbps    [3] 155.12Mbps
        [4] 155.12Mbps    [5] 155.12Mbps    [6] 155.12Mbps    [7] 155.12Mbps
To: 10.255.41.217, Local: 10.2.3.1, Remote: 10.2.3.2
Color: 0x2 fa
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
        [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
        [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
        [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
        [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
To: 10.255.41.217, Local: 172.16.30.1, Remote: 172.16.30.2
Metric: 1
Static BW: 400kbps
Reservable BW: 400kbps
Available BW [priority] bps:
        [0] 370kbps      [1] 370kbps      [2] 370kbps      [3] 370kbps
        [4] 370kbps      [5] 370kbps      [6] 370kbps      [7] 370kbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
        [0] 370kbps      [1] 370kbps      [2] 370kbps      [3] 370kbps
        [4] 370kbps      [5] 370kbps      [6] 370kbps      [7] 370kbps
To: 10.255.41.217, Local: 10.2.3.5, Remote: 10.2.3.6
Color: 0x4 backup
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
        [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
        [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
        [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
        [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
To: 10.2.5.1-1, Local: 10.2.5.1, Remote: 0.0.0.0

```



```

Color: 0x4 backup
Metric: 1
Static BW: 100Mbps
Reservable BW: 100Mbps
Available BW [priority] bps:
    [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
    [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
        [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps

user@router0> show ted database 10.255.41.217 extensive
TED database: 0 ISIS nodes 8 INET nodes
NodeID: 10.255.41.217
Type: Rtr, Age: 473 secs, LinkIn: 6, LinkOut: 6
Protocol: OSPF(0.0.0.0)
To: 10.255.41.216, Local: 10.2.3.2, Remote: 10.2.3.1
Color: 0x2 fa
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
    [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
    [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
        [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
To: 10.255.41.216, Local: 172.16.30.2, Remote: 172.16.30.1
Metric: 1
Static BW: 400kbps
Reservable BW: 400kbps
Available BW [priority] bps:
    [0] 370kbps      [1] 370kbps      [2] 370kbps      [3] 370kbps
    [4] 370kbps      [5] 370kbps      [6] 370kbps      [7] 370kbps
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:

```



```

        [0] 370kbps      [1] 370kbps      [2] 370kbps      [3] 370kbps
        [4] 370kbps      [5] 370kbps      [6] 370kbps      [7] 370kbps
To: 10.255.41.216, Local: 10.2.3.6, Remote: 10.2.3.5
Color: 0x4 backup
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
        [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
        [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
        [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
        [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
To: 10.255.41.215, Local: 10.3.5.1, Remote: 10.3.5.2
Color: 0x2 fa
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
        [0] 155.12Mbps  [1] 155.12Mbps  [2] 155.12Mbps  [3] 155.12Mbps
        [4] 155.12Mbps  [5] 155.12Mbps  [6] 155.12Mbps  [7] 155.12Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
        [0] 155.12Mbps  [1] 155.12Mbps  [2] 155.12Mbps  [3] 155.12Mbps
        [4] 155.12Mbps  [5] 155.12Mbps  [6] 155.12Mbps  [7] 155.12Mbps
To: 10.255.41.221, Local: 10.3.6.1, Remote: 10.3.6.2
Color: 0x8 other
Metric: 1
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
        [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
        [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
        [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
        [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps

```



```

To: 10.3.4.1-1, Local: 10.3.4.1, Remote: 0.0.0.0
Color: 0x4 backup
Metric: 1
Static BW: 100Mbps
Reservable BW: 100Mbps
Available BW [priority] bps:
    [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
    [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 100Mbps    [1] 100Mbps    [2] 100Mbps    [3] 100Mbps
    [4] 100Mbps    [5] 100Mbps    [6] 100Mbps    [7] 100Mbps

user@router0> show rsvp session name e2e_lsp_r0r5 extensive
Ingress RSVP: 1 sessions
10.255.41.221
  From: 10.255.41.222, LSPstate: Up, ActiveRoute: 2
  LSPname: e2e_lsp_r0r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101584
  Resv style: 1 FF, Label in: -, Label out: 101584
  Time left: -, Since: Wed Sep  7 19:02:56 2005
  Tspec: rate 30kbps size 30kbps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 29458 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.1.2.2 (so-0/0/3.0) 15 pkts
  RESV rcvfrom: 10.1.2.2 (so-0/0/3.0) 16 pkts
Explct route: 10.1.2.2 172.16.30.2 10.3.6.2
Record route: <self> 10.1.2.2 172.16.30.2 10.3.6.2
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```


Router 1

On Router 1, verify that your LMP traffic engineering link configuration is working and that the end-to-end LSP is traversing the traffic engineering link by issuing the **show link-management** set of commands. You can also issue the **show rsvp session extensive** command to confirm that the FA-LSP is operational.

```
user@router1> show link-management
```

```
Peer name: r4 , System identifier: 10758
```

```
State: Up, Control address: 10.255.41.217
```

```
TE links:
```

```
link_r1r4
```

```
TE link name: link_r1r4, State: Up
```

```
Local identifier: 16299, Remote identifier: 0, Local address: 172.16.30.1, Remote address: 172.16.30.2,
```

```
Encoding: Packet, Switching: Packet, Minimum bandwidth: 0bps, Maximum bandwidth: 400kbps,
```

```
Total bandwidth: 400kbps, Available bandwidth: 370kbps
```

Name	State	Local ID	Remote ID	Bandwidth Used	LSP-name
fa_lsp_r1r4	Up	22642	0	400kbps	Yes e2e_lsp_r0r5

```
user@router1> show link-management te-link name link_r1r4 detail
```

```
TE link name: link_r1r4, State: Up
```

```
Local identifier: 16299, Remote identifier: 0, Local address: 172.16.30.1, Remote address: 172.16.30.2,
```

```
Encoding: Packet, Switching: Packet, Minimum bandwidth: 0bps, Maximum bandwidth: 400kbps,
```

```
Total bandwidth: 400kbps, Available bandwidth: 370kbps
```

```
Resource: fa_lsp_r1r4, Type: LSP, System identifier: 2147483683, State: Up, Local identifier: 22642,
```

```
Remote identifier: 0
```

```
Total bandwidth: 400kbps, Unallocated bandwidth: 370kbps
```

```
Traffic parameters: Encoding: Packet, Switching: Packet, Granularity: Unknown
```

```
Number of allocations: 1, In use: Yes
```

```
LSP name: e2e_lsp_r0r5, Allocated bandwidth: 30kbps
```

```
user@router1> show rsvp session name fa_lsp_r1r4 extensive
```

```
Ingress RSVP: 1 sessions
```

```
10.255.41.217
```

```
From: 10.255.41.216, LSPstate: Up, ActiveRoute: 0
```

```
LSPname: fa_lsp_r1r4, LSPpath: Primary
```



```

Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 100816
Resv style: 1 FF, Label in: -, Label out: 100816
Time left: -, Since: Wed Sep 7 19:02:33 2005
Tspec: rate 400kbps size 400kbps peak Infbps m 20 M 1500
Port number: sender 2 receiver 5933 protocol 0
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.2.4.2 (so-0/0/2.0) 28 pkts
RESV rcvfrom: 10.2.4.2 (so-0/0/2.0) 26 pkts
Explot route: 10.2.4.2 10.4.5.2 10.3.5.1
Record route: <self> 10.2.4.2 10.4.5.2 10.3.5.1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 2 sessions
Total 0 displayed, Up 0, Down 0

```

RELATED DOCUMENTATION

[RSVP LSP Tunnels Overview | 885](#)

[Configuring Link Management Protocol Peers | 912](#)

[Configuring Link Management Protocol Traffic Engineering Links | 913](#)

[Configuring Peer Interfaces in OSPF and RSVP | 914](#)

[Defining Label-Switched Paths for the FA-LSP | 915](#)

[Establishing FA-LSP Path Information | 915](#)

[Option: Tearing Down RSVP LSPs Gracefully | 916](#)

Configuring Link Management Protocol Peers

After you set up traffic engineering links, configure LMP network peers with the **peer *peer-name*** statement at the **[edit protocols link-management]** hierarchy level. A peer is the network device with which your routing platform communicates and establishes an FA-LSP. Designate a peer name, configure the peer router ID as the address (often a loopback address), and apply the traffic engineering link to be associated

with this peer. Remember to configure both sides of a peering relationship to enable bidirectional communication.

Unlike GMPLS, you must not configure a control channel for a peer. If you include a control channel, the commit operation fails.

```
[edit]
protocols {
  link-management {
    peer peer-name { # Configure the name of your network peer.
      address ip-address; # Include the router ID of the peer.
      te-link te-link-name; # Assign a TE link to this peer.
    }
  }
}
```

RELATED DOCUMENTATION

[Example: RSVP LSP Tunnel Configuration](#) | 888

Configuring Link Management Protocol Traffic Engineering Links

To begin your RSVP LSP tunnel configuration, configure LMP traffic engineering links on both the ingress and egress routing platforms. Because traffic engineering links define a unidirectional connection between peer devices, you must configure traffic engineering links in both directions between peers to enable the bidirectional transport of packets.

To configure traffic engineering links in LMP, include the **te-link *te-link-name*** statement at the **[edit protocols link-management]** hierarchy level. Define the traffic engineering link options shown below, especially the label-switched path to be used as the FA-LSP to reach the peer. Optionally, you can specify the traffic engineering metric for the traffic engineering link (TE link). By default, the traffic engineering metric is derived from the CSPF computation.

```
[edit]
protocols {
  link-management {
    te-link te-link-name { # Name of the TE link.
      label-switched-path lsp-name; # LSP used for the forwarding adjacency.
      local-address ip-address; # Local IP address associated with the TE link.
      remote-address ip-address; # Remote IP address mapped to the TE link.
    }
  }
}
```



```

        te-metric value; # Traffic engineering metric used for the TE link.
    }
}
}

```

RELATED DOCUMENTATION

[Example: RSVP LSP Tunnel Configuration](#) | 888

Configuring Peer Interfaces in OSPF and RSVP

After you establish LMP peers, you must add peer interfaces to OSPF and RSVP. A peer interface is a virtual interface used to support the control adjacency between two peers.

The peer interface name must match the **peer *peer-name*** statement configured in LMP at the **[edit protocols link-management]** hierarchy level. Because actual protocol packets are sent and received by peer interfaces, the peer interfaces can be signaled and advertised to peers like any other physical interface configured for OSPF and RSVP. To configure OSPF routing for LMP peers, include the **peer-interface** statement at the **[edit protocols ospf area *area-number*]** hierarchy level. To configure RSVP signaling for LMP peers, include the **peer-interface** statement at the **[edit protocols rsvp]** hierarchy level.

```

[edit]
protocols {
  rsvp {
    peer-interface peer-name { # Configure the name of your LMP peer.
    }
  }
  ospf {
    area area-number {
      peer-interface peer-name { # Configure the name of your LMP peer.
      }
    }
  }
}
}

```

RELATED DOCUMENTATION

[Example: RSVP LSP Tunnel Configuration](#) | 888

Defining Label-Switched Paths for the FA-LSP

Next, define your FA-LSP by including the **label-switched-path** statement at the **[edit protocols mpls]** hierarchy level. Include the router ID of the peer in the **to** statement at the **[edit protocols mpls label-switched-path]** hierarchy level. Because packet LSPs are unidirectional, you must create one FA-LSP to reach the peer and a second FA-LSP to return from the peer.

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      from ip-address;
      to ip-address;
      primary path-name;
      secondary path-name;
      no-cspf; # This statement to disable CSPF is optional.
    }
  }
}
```

RELATED DOCUMENTATION

[Example: RSVP LSP Tunnel Configuration | 888](#)

Establishing FA-LSP Path Information

When you configure explicit LSP paths for an FA-LSP, you must use the traffic engineering link remote address as your next-hop address. When CSPF is supported, you can use any path option you wish. However, when CSPF is disabled with the **no-cspf** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level, you must use strict paths.

```
[edit]
protocols {
  mpls {
    path path-name {
      next-hop-address (strict | loose);
    }
  }
}
```



```
}
```

NOTE: If the end-to-end LSP originates on the same routing platform as the FA-LSP, you must disable CSPF and use strict paths.

RELATED DOCUMENTATION

[Example: RSVP LSP Tunnel Configuration | 888](#)

Option: Tearing Down RSVP LSPs Gracefully

You can tear down an RSVP LSP in a two-step process that gracefully withdraws the RSVP session used by the LSP. For all neighbors that support graceful teardown, a request for the teardown is sent by the routing platform to the destination endpoint for the LSP and all RSVP neighbors in the path. The request is included within the **ADMIN_STATUS** field of the RSVP packet. When neighbors receive the request, they prepare for the RSVP session to be withdrawn. A second message is sent by the routing platform to complete the teardown of the RSVP session. If a neighbor does not support graceful teardown, the request is handled as a standard session teardown rather than a graceful one.

To perform a graceful teardown of an RSVP session, issue the **clear rsvp session gracefully** command. Optionally, you can specify the source and destination address of the RSVP session, the LSP identifier of the RSVP sender, and the tunnel identifier of the RSVP session. To use these qualifiers, include the **connection-source**, **connection-destination**, **lsp-id**, and **tunnel-id** options when you issue the **clear rsvp session gracefully** command.

You can also configure the amount of time that the routing platform waits for neighbors to receive the graceful teardown request before initiating the actual teardown by including the **graceful-deletion-timeout** statement at the **[edit protocols rsvp]** hierarchy level. The default graceful deletion timeout value is 30 seconds, with a minimum value of 1 second and a maximum value of 300 seconds. To view the current value configured for graceful deletion timeout, issue the **show rsvp version** operational mode command.

RELATED DOCUMENTATION

[Example: RSVP LSP Tunnel Configuration | 888](#)

Configuring LDP

IN THIS CHAPTER

- LDP Introduction | 918
- Understanding the LDP Signaling Protocol | 919
- Example: Configuring LDP-Signaled LSPs | 920
- Supported LDP Standards | 923
- Junos OS LDP Protocol Implementation | 924
- LDP Operation | 924
- LDP Message Types | 925
- Discovery Messages | 927
- Session Messages | 927
- Advertisement Messages | 927
- Notification Messages | 927
- Tunneling LDP LSPs in RSVP LSPs | 928
- Tunneling LDP LSPs in RSVP LSPs Overview | 928
- Label Operations | 930
- LDP Session Protection | 931
- LDP Native IPv6 Support Overview | 932
- Longest Match Support for LDP Overview | 933
- Minimum LDP Configuration | 933
- Enabling and Disabling LDP | 934
- Configuring the LDP Timer for Hello Messages | 934
- Configuring the Delay Before LDP Neighbors Are Considered Down | 935
- Enabling Strict Targeted Hello Messages for LDP | 937
- Configuring the Interval for LDP Keepalive Messages | 938
- Configuring the LDP Keepalive Timeout | 938
- Configuring Longest Match for LDP | 939
- Example: Configuring Longest Match for LDP | 940
- Configuring LDP Route Preferences | 959
- LDP Graceful Restart | 959

- [Configuring LDP Graceful Restart | 960](#)
- [Filtering Inbound LDP Label Bindings | 963](#)
- [Filtering Outbound LDP Label Bindings | 965](#)
- [Specifying the Transport Address Used by LDP | 968](#)
- [Control Transport Address Used for Targeted-LDP Session | 969](#)
- [Configuring the Prefixes Advertised into LDP from the Routing Table | 971](#)
- [Configuring FEC Deaggregation | 973](#)
- [Configuring Policers for LDP FECs | 974](#)
- [Configuring LDP IPv4 FEC Filtering | 975](#)
- [Configuring BFD for LDP LSPs | 975](#)
- [Configuring ECMP-Aware BFD for LDP LSPs | 979](#)
- [Configuring a Failure Action for the BFD Session on an LDP LSP | 980](#)
- [Configuring the Holddown Interval for the BFD Session | 980](#)
- [Configuring LDP Link Protection | 981](#)
- [Example: Configuring LDP Link Protection | 983](#)
- [Understanding Multicast-Only Fast Reroute | 1014](#)
- [Configuring Multicast-Only Fast Reroute | 1021](#)
- [Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)
- [Example: Configuring LDP Downstream on Demand | 1046](#)
- [Configuring LDP Native IPv6 Support | 1053](#)
- [Example: Configuring LDP Native IPv6 Support | 1054](#)
- [Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs | 1072](#)
- [LDP Mapping Server for Interoperability of Segment Routing with LDP Overview | 1110](#)
- [Configuring Miscellaneous LDP Properties | 1115](#)
- [Configuring LDP LSP Traceroute | 1123](#)
- [Collecting LDP Statistics | 1124](#)
- [Tracing LDP Protocol Traffic | 1127](#)

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.

Understanding the LDP Signaling Protocol

LDP is a signaling protocol that runs on a device configured for MPLS support. The successful configuration of both MPLS and LDP initiates the exchange of TCP packets across the LDP interfaces. The packets establish TCP-based LDP sessions for the exchange of MPLS information within the network. Enabling both MPLS and LDP on the appropriate interfaces is sufficient to establish LSPs.

LDP is a simple, fast-acting signaling protocol that automatically establishes LSP adjacencies within an MPLS network. Routers then share LSP updates such as hello packets and LSP advertisements across the adjacencies. Because LDP runs on top of an IGP such as IS-IS or OSPF, you must configure LDP and the IGP on the same set of interfaces. After both are configured, LDP begins transmitting and receiving LDP messages through all LDP-enabled interfaces. Because of LDP's simplicity, it cannot perform the true traffic engineering which RSVP can perform. LDP does not support bandwidth reservation or traffic constraints.

When you configure LDP on a label-switching router (LSR), the router begins sending LDP discovery messages out all LDP-enabled interfaces. When an adjacent LSR receives LDP discovery messages, it establishes an underlying TCP session. An LDP session is then created on top of the TCP session. The TCP three-way handshake ensures that the LDP session has bidirectional connectivity. After they establish the LDP session, the LDP neighbors maintain, and terminate, the session by exchanging messages. LDP advertisement messages allow LSRs to exchange label information to determine the next hops within a particular LSP. Any topology changes, such as a router failure, generate LDP notifications that can terminate the LDP session or generate additional LDP advertisements to propagate an LSP change.

RELATED DOCUMENTATION

[MPLS Traffic Engineering and Signaling Protocols Overview | 1134](#)

[Example: Configuring LDP-Signaled LSPs | 920](#)

Example: Configuring LDP-Signaled LSPs

IN THIS SECTION

- Requirements | 920
- Overview | 920
- Configuration | 921
- Verification | 923

This example shows how to create and configure LDP instances within an MPLS network.

Requirements

Before you begin:

- Configure network interfaces. See *Interfaces User Guide for Security Devices*.
- Configure an IGP across your network. (The LDP configuration is added to the existing IGP configuration and included in the MPLS configuration.)
- Configure a network to use LDP for LSP establishment by enabling MPLS on all transit interfaces in the MPLS network.

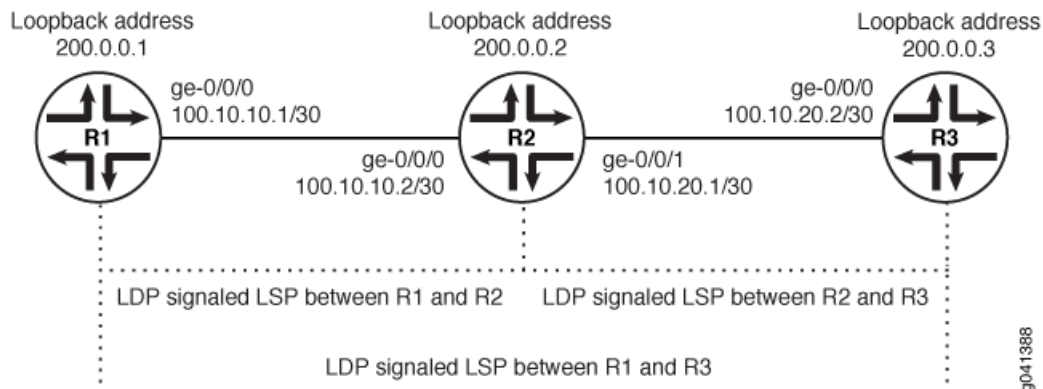
NOTE: Because LDP runs on top of an IGP such as IS-IS or OSPF, you must configure LDP and the IGP on the same set of interfaces.

Overview

To configure LDP-signaled LSPs, you must enable the MPLS family on all transit interfaces in the MPLS network and include all the transit interfaces under the **[protocols mpls]** and **[protocols ldp]** hierarchy levels.

In this example, you enable the MPLS family and create an LDP instance on all the transit interfaces. Additionally, you enable the MPLS process on all the transit interfaces in the MPLS network. In this example, you configure a sample network as shown in [Figure 70 on page 921](#).

Figure 70: Typical LDP-Signaled LSP



Configuration

CLI Quick Configuration

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

For Router R1, perform the following:

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols mpls ge-0/0/0 unit 0
set protocols ldp interface ge-0/0/0.0 unit 0
```

For Router R2, perform the following:

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols mpls ge-0/0/0 unit 0
set protocols ldp interface ge-0/0/0.0 unit 0
set interfaces ge-0/0/1 unit 0 family mpls
set protocols mpls ge-0/0/1 unit 0
set protocols ldp interface ge-0/0/0.1 unit 0
```

For Router R3, perform the following:

```
set interfaces ge-0/0/0 unit 0 family mpls
set protocols mpls ge-0/0/0 unit 0
set protocols ldp interface ge-0/0/0.0 unit 0
```


Step-by-Step Procedure

To enable LDP instances within an MPLS network:

1. Enable the MPLS family on the transit interface on Router R1.

```
[edit]
user@R1# set interfaces ge-0/0/0 unit 0 family mpls
```

2. Enable the MPLS process on the transit interface.

```
[edit]
user@R1# set protocols mpls interface ge-0/0/0 unit 0
```

3. Create the LDP instance on the transit interface.

```
[edit]
user@R1# set protocols ldp interface ge-0/0/0 unit 0
```

Results

Confirm your configuration by entering the **show** command from configuration mode. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
user@R1# show
...
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 10.100.37.20/24;
      }
      family mpls;
    }
  }
}
...
protocols {
  mpls {
    interface all;
```



```

    }
    ldp {
        interface ge-0/0/0.0;
    }
}

```

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

Verification

Confirm that the configuration is working properly.

RELATED DOCUMENTATION

[Understanding the LDP Signaling Protocol | 919](#)

Supported LDP Standards

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

- RFC 3212, *Constraint-Based LSP Setup using LDP*
- RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol*
- Internet draft draft-napierala-mpls-targeted-mlldp-01.txt, *Using LDP Multipoint Extensions on Targeted LDP Sessions*

The following RFCs do not define standards, but provide information about LDP. The IETF classifies them as “Informational.”

- RFC 3215, *LDP State Machine*
- RFC 5036, *LDP Specification*

For the following features described in the indicated sections of the RFC, Junos OS supports one of the possible modes but not the others:

- Label distribution control (section 2.6.1): Ordered mode is supported, but not Independent mode.
- Label retention (section 2.6.2): Liberal mode is supported, but not Conservative mode.
- Label advertisement (section 2.6.3): Both Downstream Unsolicited mode and Downstream on Demand mode are supported.

- RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*
- RFC 5443, *LDP IGP Synchronization*
- RFC 5561, *LDP Capabilities*
- RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

Junos OS support limited to point-to-multipoint extensions for LDP.

RELATED DOCUMENTATION

[Supported GMPLS Standards | 1356](#)

[Supported MPLS Standards | 20](#)

[Supported RSVP Standards | 816](#)

Accessing Standards Documents on the Internet

Junos OS LDP Protocol Implementation

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

LDP Operation

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

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP

session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

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

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

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

RELATED DOCUMENTATION

| [Logical Interfaces](#)

LDP Message Types

IN THIS SECTION

- [Discovery Messages | 925](#)
- [Session Messages | 926](#)
- [Advertisement Messages | 926](#)
- [Notification Messages | 926](#)

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

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.

Discovery Messages

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

LDP uses the following discovery procedures:

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

Session Messages

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

Advertisement Messages

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

Notification Messages

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

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

Tunneling LDP LSPs in RSVP LSPs

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

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 928](#)
- [Label Operations on page 930](#)

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.

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 forwarding equivalence classes (FECs). Each domain uses intra-domain LDP-over-RSVP LSP for MPLS forwarding.

NOTE: With the introduction of the multi-instance support for LDP-over-RSVP LSPs, you cannot enable MPLS on an interface that is already assigned to another routing instance. Adding an interface that is part of another routing instance at the **[edit protocols mpls]** hierarchy level, throws a configuration error at the time of commit.

Benefits of Tunneling LDP LSPs in RSVP LSPs

Tunneling LDP LSPs in RSVP LSPs provides the following benefits:

- Provides convergence of different traffic types such as IPv4, IPv6, unicast, and multicast across Layer 2 and Layer 3 VPNs.
- Enables flexible access connectivity options that can accommodate multiple topologies, different protocols, and multiple administrative boundaries.
- Enables secure interworking among multiple providers.
- Enables provision of differentiated services on a per customer basis because RSVP-TE supports traffic engineering, bandwidth guarantees, and link and node redundancy capabilities.
- Reduces the number of LSPs required in the core, which reduces the resource requirements of the protocols and routers as well as reducing convergence time.
- Provides cost-efficient rollouts with minimal network disruption because the LSPs are built using point-to-point TE tunnels to directly attached neighbors. These TE tunnels only go to the next hop, not end to end. Then when LDP is run over those tunnels, the sessions are built to the directly connected neighbor. When there is a change in the network, such as adding a new node, the directly connected neighbors of the new node have RSVP and LDP sessions. Thus, the RSVP LSPs are only to the next hop, and LDP takes care of advertising labels for the new addresses.

Release History Table

Release	Description
15.1	Beginning with Junos OS Release 15.1, multi-instance support is extended to LDP over RSVP tunneling for a virtual router routing instance.

RELATED DOCUMENTATION

[Label Operations](#) | 930

Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP

Label Operations

Figure 71 on page 930 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see “MPLS Label Overview” on page 419.) 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 71: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs





Figure 72 on page 931 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 72: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs



LDP Session Protection

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

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

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

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

LDP Native IPv6 Support Overview

IPv6 connectivity often relies on tunneling IPv6 over an IPv4 MPLS core with IPv4-signaled MPLS label-switched paths (LSPs). This requires the IPv4-signaled LSPs to be configured statically or established dynamically by IPv6 provider edge routers. Because of the growing demand of IPv6, it has become imperative to deploy an IPv6 MPLS core with an IPv6-signaled LSP to provide IPv6 connectivity. In Junos OS, LDP is supported in an IPv6 network only, and in an IPv6/IPv4 dual-stack network as described in *RFC 7552*. Apart from providing a single session for both IPv4 and IPv6 networks, Junos OS LDP supports separate IPv4 sessions for IPv4 only, and IPv6 sessions for IPv6 only.

You can configure the address family as **inet** for IPv4 or **inet6** for IPv6, or both. If the family address is not configured, then the default address of family inet is enabled. When both IPv4 and IPv6 are configured, you can use the **transport-preference** statement to configure the preferred transport to be either **IPv4** or **IPv6**. Based on the preference, LDP attempts to establish a TCP connection using IPv4 or IPv6. By default, IPv6 is selected. The **dual-transport** statement allows Junos OS LDP to establish the TCP connection over IPv4 with IPv4 neighbors, and over IPv6 with IPv6 neighbors as a single-stack LSR. The **inet-lsr-id** and **inet6-lsr-id** IDs are the two LSR IDs that have to be configured to establish an LDP session over IPv4 and IPv6 TCP transport. These two IDs should be non-zero and must be configured with different values.

RELATED DOCUMENTATION

[Example: Configuring LDP Native IPv6 Support | 1054](#)

[Configuring LDP Native IPv6 Support | 1053](#)

[dual-transport | 2592](#)

[family \(Protocols LDP\) | 2614](#)

Longest Match Support for LDP Overview

LDP is often used to establish MPLS label-switched paths (LSPs) throughout a complete network domain using an IGP such as OSPF or IS-IS. In such a network, all links in the domain have IGP adjacencies as well as LDP adjacencies. LDP establishes the LSPs on the shortest path to a destination as determined by IGP. In Junos OS, the LDP implementation does an exact match lookup on the IP address of the forwarding equivalence class (FEC) in the routing information base (RIB) or IGP routes for label mapping. This exact mapping requires MPLS end-to-end LDP endpoint IP addresses to be configured in all the label edge routers (LERs). This defeats the purpose of IP hierarchical design or default routing in access devices. Configuring **longest-match** allows LDP to set up LSP based on the routes aggregated or summarized across OSPF areas or IS-IS levels in the inter-domain.

RELATED DOCUMENTATION

[longest-match](#) | 2667

[Example: Configuring Longest Match for LDP](#) | 940

[Configuring Longest Match for LDP](#) | 939

Minimum LDP Configuration

To enable LDP with minimal configuration:

1. Enable all relevant interfaces under family MPLS. In the case of directed LDP, the loopback interface needs to be enabled with family MPLS.
2. (Optional) Configure the relevant interfaces under the **[edit protocol mpls]** hierarchy level.
3. 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.

RELATED DOCUMENTATION

| [ldp](#) | [2944](#)

Enabling and Disabling LDP

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

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

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

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

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

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

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

Configuring the LDP Timer for Hello Messages

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

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.

- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 935](#).

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 934](#).

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an LDP peer router advertises a shorter hold time than the value you have configured, the peer router’s advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

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

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

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

Configuring the LDP Hold Time for Link Hello Messages

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

```
hold-time seconds;
```

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

Configuring the LDP Hold Time for Targeted Hello Messages

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

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

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

Enabling Strict Targeted Hello Messages for LDP

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

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

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

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

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

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

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

Configuring the Interval for LDP Keepalive Messages

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

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 935](#).

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 Longest Match for LDP

In order to allow LDP to learn the routes aggregated or summarized across OSPF areas or ISIS levels in inter-domain, Junos OS allows you to configure longest match for LDP based on *RFC5283*.

Before you configure longest match for LDP, you must do the following:

1. Configure the device interfaces.
2. Configure the MPLS protocol.
3. Configure the OSPF protocol.

To configure longest match for LDP, you must do the following:

1. Configure longest match for the LDP protocol.

```
[edit protocols ldp]
user@host# set longest-match
```

2. Configure the LDP protocol on the interface.

```
[edit protocols ldp]
user@host# set interface interface-name
```

For example, to configure the interfaces:

```
[edit protocols ldp]
user@host# set interface ge-0/0/2.0
user@host# set interface lo0.0
```

RELATED DOCUMENTATION

[Example: Configuring Longest Match for LDP | 940](#)

[Longest Match Support for LDP Overview | 933](#)

[longest-match | 2667](#)

Example: Configuring Longest Match for LDP

IN THIS SECTION

- Requirements | 940
- Overview | 940
- Configuration | 941
- Verification | 948

This example shows how to configure longest match for LDP based on *RFC5283*. This allows LDP to learn the routes aggregated or summarized across OSPF areas or ISIS levels in inter-domain.. The longest match policy provides per prefix granularity.

Requirements

This example uses the following hardware and software components:

- Six MX Series routers with OSPF protocol, and LDP enabled on the connected interfaces.
- Junos OS Release 16.1 or later running on all devices.

Before you begin:

- Configure the device interfaces.
- Configure OSPF.


Overview

LDP is often used to establish MPLS label-switched paths (LSPs) throughout a complete network domain using an IGP such as OSPF or IS-IS. In such a network, all links in the domain have IGP adjacencies as well as LDP adjacencies. LDP establishes the LSPs on the shortest path to a destination as determined by IP forwarding. In Junos OS, the LDP implementation does an exact match lookup on the IP address of the FEC in the RIB or IGP routes for label mapping. This exact mapping requires MPLS end-to-end LDP endpoint IP addresses to be configured in all the LERs. This defeats the purpose of IP hierarchical design or default routing in access devices. Configuring **longest-match** helps to overcome this by suppressing the exact match behaviour and setup LSP based on the longest matching route on per-prefix basis.

Topology

In the topology, [Figure 73 on page 941](#) shows the longest match for LDP is configured on Device R0 .

Figure 73: Example Longest Match for LDP



Configuration

CLI Quick Configuration

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

R0

```
set interfaces ge-0/0/0 unit 0 family inet address 22.22.22.1/24
set interfaces ge-0/0/1 unit 0 family inet address 15.15.15.1/24
set interfaces ge-0/0/2 unit 0 family inet address 11.11.11.1/24
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.112.1/32 primary
set interfaces lo0 unit 0 family inet address 10.255.112.1/32 preferred
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0001.00
set routing-options router-id 10.255.112.1
set protocols mpls interface ge-0/0/2.0
set protocols ospf area 0.0.0.1 interface ge-0/0/2.0
set protocols ospf area 0.0.0.1 interface lo0.0 passive
set protocols ldp longest-match
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface lo0.0
```

R1


```

set interfaces ge-0/0/0 unit 0 family inet address 11.11.11.2/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 12.12.12.1/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.112.2/32 primary
set interfaces lo0 unit 0 family inet address 10.255.112.2/32 preferred
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0002.00
set routing-options router-id 10.255.112.2
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.1 interface ge-0/0/0.0
set protocols ldp longest-match
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0

```

R2

```

set interfaces ge-0/0/0 unit 0 family inet address 24.24.24.1/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 12.12.12.2/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 23.23.23.1/24
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 22.22.22.2/24
set interfaces ge-0/0/4 unit 0 family inet address 25.25.25.1/24
set interfaces ge-0/0/4 unit 0 family iso
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.111.4/32 primary
set interfaces lo0 unit 0 family inet address 10.255.111.4/32 preferred
set interfaces lo0 unit 0 family iso address 49.0003.0192.0168.0003.00
set routing-options router-id 10.255.111.4
set protocols mpls interface ge-0/0/1.0

```



```

set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/4.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.2 area-range 10.255.111.0/24
set protocols ospf area 0.0.0.2 interface ge-0/0/2.0
set protocols ospf area 0.0.0.2 interface ge-0/0/0.0
set protocols ospf area 0.0.0.2 interface ge-0/0/4.0
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface ge-0/0/2.0
set protocols ldp interface ge-0/0/4.0
set protocols ldp interface lo0.0

```

R3

```

set interfaces ge-0/0/0 unit 0 family inet address 35.35.35.1/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 23.23.23.2/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 34.34.34.1/24
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.111.1/32 primary
set interfaces lo0 unit 0 family inet address 10.255.111.1/32 preferred
set interfaces lo0 unit 0 family iso address 49.0003.0192.0168.0004.00
set routing-options router-id 10.255.111.1
set protocols mpls interface ge-0/0/1.0
set protocols ospf area 0.0.0.2 interface ge-0/0/1.0
set protocols ospf area 0.0.0.2 interface fxp0.0 disable
set protocols ospf area 0.0.0.2 interface lo0.0 passive
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0

```

R4


```

set interfaces ge-0/0/0 unit 0 family inet address 45.45.45.1/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 24.24.24.2/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 34.34.34.2/24
set interfaces ge-0/0/2 unit 0 family iso
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.111.2/32 primary
set interfaces lo0 unit 0 family inet address 10.255.111.2/32 preferred
set interfaces lo0 unit 0 family iso address 49.0003.0192.0168.0005.00
set routing-options router-id 10.255.111.2
set protocols mpls interface ge-0/0/1.0
set protocols ospf area 0.0.0.2 interface ge-0/0/1.0
set protocols ospf area 0.0.0.2 interface fxp0.0 disable
set protocols ospf area 0.0.0.2 interface lo0.0 passive
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0

```

R5

```

set interfaces ge-0/0/0 unit 0 family inet address 25.25.25.2/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 15.15.15.2/24
set interfaces ge-0/0/2 unit 0 family inet address 35.35.35.2/24
set interfaces ge-0/0/3 unit 0 family inet address 45.45.45.2/24
set interfaces ge-0/0/3 unit 0 family iso
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.111.3/32 primary
set interfaces lo0 unit 0 family inet address 10.255.111.3/32 preferred
set interfaces lo0 unit 0 family iso address 49.0003.0192.0168.0006.00
set routing-options router-id 10.255.111.3
set protocols mpls interface ge-0/0/0.0
set protocols ospf area 0.0.0.2 interface ge-0/0/0.0
set protocols ospf area 0.0.0.2 interface fxp0.0 disable
set protocols ospf area 0.0.0.2 interface lo0.0 passive
set protocols ldp interface ge-0/0/0.0
set protocols ldp interface lo0.0

```


Configuring Device R0

Step-by-Step Procedure

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

To configure Device R0:

1. Configure the interfaces.

```
[edit interfaces]
set ge-0/0/0 unit 0 family inet address 22.22.22.1/24

set ge-0/0/1 unit 0 family inet address 15.15.15.1/24

set ge-0/0/2 unit 0 family inet address 11.11.11.1/24
set ge-0/0/2 unit 0 family iso
set ge-0/0/2 unit 0 family mpls
```

2. Assign the loopback addresses to the device.

```
[edit interfaces lo0 unit 0 family]
set inet address 10.255.112.1/32 primary
set inet address 10.255.112.1/32 preferred
set iso address 49.0002.0192.0168.0001.00
```

3. Configure the router ID.

```
[edit routing-options]
set router-id 10.255.112.1
```

4. Configure the MPLS protocol on the interface.

```
[edit protocols mpls]
set interface ge-0/0/2.0
```

5. Configure the OSPF protocol on the interface.

```
[edit protocols ospf]
set area 0.0.0.1 interface ge-0/0/2.0
```



```
set area 0.0.0.1 interface lo0.0 passive
```

6. Configure longest match for the LDP protocol.

```
[edit protocols ldp]
set longest-match
```

7. Configure the LDP protocol on the interface.

```
[edit protocols ldp]
set interface ge-0/0/2.0
set interface lo0.0
```

Results

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

```
user@R0# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 22.22.22.1/24;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 15.15.15.1/24;
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 11.11.11.1/24;
    }
    family iso;
```



```

        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.112.1/32 {
                primary;
                preferred;
            }
        }
        family iso {
            address 49.0002.0192.0168.0001.00;
        }
    }
}
}

```

```

user@R0# show protocols
mpls {
    interface ge-0/0/2.0;
}
ospf {
    area 0.0.0.1 {
        interface ge-0/0/2.0;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    longest-match;
    interface ge-0/0/2.0;
    interface lo0.0;
}

```

```

user@R0# show routing-options
router-id 10.255.112.1;

```

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

Verification

IN THIS SECTION

- [Verifying the Routes | 948](#)
- [Verifying LDP Overview Information | 953](#)
- [Verify the LDP Entries in the Internal Topology Table | 954](#)
- [Verify Only FEC Information of LDP Route | 956](#)
- [Verify FEC and Shadow Routes of LDP | 957](#)

Confirm that the configuration is working properly.

Verifying the Routes

Purpose

Verify that the expected routes are learned.

Action

On Device R0, from operational mode, run the **show route** command to display the routes in the routing table.

```
user@R0> show route
```

```
inet.0: 62 destinations, 62 routes (62 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.4.0.0/16      *[Static/5] 10:08:01
                  > to 10.92.31.254 via fxp0.0
10.5.0.0/16      *[Static/5] 10:08:01
                  > to 10.92.31.254 via fxp0.0
10.6.128.0/17    *[Static/5] 10:08:01
                  > to 10.92.31.254 via fxp0.0
10.9.0.0/16      *[Static/5] 10:08:01
                  > to 10.92.31.254 via fxp0.0
10.10.0.0/16     *[Static/5] 10:08:01
                  > to 10.92.31.254 via fxp0.0
10.13.4.0/23     *[Static/5] 10:08:01
                  > to 10.92.31.254 via fxp0.0
```



```

10.13.10.0/23      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.82.0.0/15       *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.84.0.0/16       *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.85.12.0/22      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.92.0.0/16       *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.92.16.0/20      *[Direct/0] 10:08:01
                   > via fxp0.0
10.92.20.175/32    *[Local/0] 10:08:01
                   Local via fxp0.0
10.94.0.0/16       *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.99.0.0/16       *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.102.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.150.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.155.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.157.64.0/19     *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.160.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.204.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.205.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.206.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.207.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.209.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.212.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.213.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.214.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0

```



```

10.215.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.216.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.218.13.0/24     *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.218.14.0/24     *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.218.16.0/20     *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.218.32.0/20     *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.227.0.0/16      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
10.255.111.0/24    *[OSPF/10] 09:52:14, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0
10.255.111.4/32    *[OSPF/10] 09:54:10, metric 2
                   > to 11.11.11.2 via ge-0/0/2.0
10.255.112.1/32    *[Direct/0] 09:55:05
                   > via lo0.0
10.255.112.2/32    *[OSPF/10] 09:54:18, metric 1
                   > to 11.11.11.2 via ge-0/0/2.0
11.11.11.0/24      *[Direct/0] 09:55:05
                   > via ge-0/0/2.0
11.11.11.1/32      *[Local/0] 09:55:05
                   Local via ge-0/0/2.0
12.12.12.0/24      *[OSPF/10] 09:54:18, metric 2
                   > to 11.11.11.2 via ge-0/0/2.0
15.15.15.0/24      *[Direct/0] 09:55:05
                   > via ge-0/0/1.0
15.15.15.1/32      *[Local/0] 09:55:05
                   Local via ge-0/0/1.0
22.22.22.0/24      *[Direct/0] 09:55:05
                   > via ge-0/0/0.0
22.22.22.1/32      *[Local/0] 09:55:05
                   Local via ge-0/0/0.0
23.23.23.0/24      *[OSPF/10] 09:54:10, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0
24.24.24.0/24      *[OSPF/10] 09:54:10, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0
25.25.25.0/24      *[OSPF/10] 09:54:10, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0
128.92.17.45/32    *[OSPF/10] 09:54:05, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0

```



```

128.92.20.175/32    *[Direct/0] 10:08:01
                   > via lo0.0
128.92.21.186/32    *[OSPF/10] 09:54:10, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0
128.92.25.135/32    *[OSPF/10] 09:54:10, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0
128.92.27.91/32     *[OSPF/10] 09:54:18, metric 1
                   > to 11.11.11.2 via ge-0/0/2.0
128.92.28.70/32     *[OSPF/10] 09:54:10, metric 2
                   > to 11.11.11.2 via ge-0/0/2.0
172.16.0.0/12       *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
192.168.0.0/16       *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
192.168.102.0/23     *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
207.17.136.0/24      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
207.17.136.192/32    *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
207.17.137.0/24      *[Static/5] 10:08:01
                   > to 10.92.31.254 via fxp0.0
224.0.0.5/32         *[OSPF/10] 09:55:05, metric 1
                   MultiRecv

inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.111.1/32      *[LDP/9] 09:41:03, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0, Push 300128
10.255.111.2/32      *[LDP/9] 09:41:03, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0, Push 300144
10.255.111.3/32      *[LDP/9] 09:41:03, metric 3
                   > to 11.11.11.2 via ge-0/0/2.0, Push 300160
10.255.111.4/32      *[LDP/9] 09:54:10, metric 2, tag 0
                   > to 11.11.11.2 via ge-0/0/2.0, Push 300000
10.255.112.2/32      *[LDP/9] 09:54:48, metric 1, tag 0
                   > to 11.11.11.2 via ge-0/0/2.0

iso.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1280.9202.0175/152
                   *[Direct/0] 10:08:01

```



```

                > via lo0.0
49.0002.0192.0168.0001/72
                *[Direct/0] 09:55:05
                > via lo0.0

mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 09:55:05, metric 1
                Receive
1                *[MPLS/0] 09:55:05, metric 1
                Receive
2                *[MPLS/0] 09:55:05, metric 1
                Receive
13               *[MPLS/0] 09:55:05, metric 1
                Receive
300064            *[LDP/9] 09:54:48, metric 1
                > to 11.11.11.2 via ge-0/0/2.0, Pop
300064(S=0)       *[LDP/9] 09:54:48, metric 1
                > to 11.11.11.2 via ge-0/0/2.0, Pop
300112            *[LDP/9] 09:54:10, metric 2, tag 0
                > to 11.11.11.2 via ge-0/0/2.0, Swap 300000
300192            *[LDP/9] 09:41:03, metric 3
                > to 11.11.11.2 via ge-0/0/2.0, Swap 300128
300208            *[LDP/9] 09:41:03, metric 3
                > to 11.11.11.2 via ge-0/0/2.0, Swap 300144
300224            *[LDP/9] 09:41:03, metric 3
                > to 11.11.11.2 via ge-0/0/2.0, Swap 300160

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

abcd::128:92:20:175/128
                *[Direct/0] 10:08:01
                > via lo0.0
fe80::5668:a50f:fcc1:1f9c/128
                *[Direct/0] 10:08:01
                > via lo0.0

```

Meaning

The output shows all the routes in the routing table of Device R0.

Verifying LDP Overview Information

Purpose

Display LDP overview information.

Action

On Device R0, from operational mode, run the **show ldp overview** command to display the overview of the LDP.

user@R0> **show ldp overview**

```
Instance: master
  Reference count: 2
  Router ID: 10.255.112.1
  Message id: 8
  Configuration sequence: 6
  Deaggregate: disabled
  Explicit null: disabled
  IPv6 tunneling: disabled
  Strict targeted hellos: disabled
  Loopback if added: yes
  Route preference: 9
  Unicast transit LSP chaining: disabled
  P2MP transit LSP chaining: disabled
  Transit LSP statistics based on route statistics: disabled
  LDP route acknowledgement: enabled
  LDP mtu discovery: disabled
  Longest Match: enabled
  Capabilities enabled: none
  Egress FEC capabilities enabled: entropy-label-capability
  Downstream unsolicited Sessions:
    Operational: 1
    Retention: liberal
    Control: ordered
  Auto targeted sessions:
    Auto targeted: disabled
  Timers:
    Keepalive interval: 10, Keepalive timeout: 30
    Link hello interval: 5, Link hello hold time: 15
    Targeted hello interval: 15, Targeted hello hold time: 45
    Label withdraw delay: 60, Make before break timeout: 30
    Make before break switchover delay: 3
    Link protection timeout: 120
  Graceful restart:
```



```

Restart: disabled, Helper: enabled, Restart in process: false
Reconnect time: 60000, Max neighbor reconnect time: 120000
Recovery time: 160000, Max neighbor recovery time: 240000
Traffic Engineering:
  Bgp igp: disabled
  Both ribs: disabled
  Mpls forwarding: disabled
IGP:
  Tracking igp metric: disabled
  Sync session up delay: 10
Session protection:
  Session protection: disabled
  Session protection timeout: 0
Interface addresses advertising:
  11.11.11.1
  10.255.112.1
  128.92.20.175
Label allocation:
  Current number of LDP labels allocated: 5
  Total number of LDP labels allocated: 11
  Total number of LDP labels freed: 6
  Total number of LDP label allocation failure: 0
  Current number of labels allocated by all protocols: 5

```

Meaning

The output displays the LDP overview information of Device R0

Verify the LDP Entries in the Internal Topology Table

Purpose

Display the route entries in the Label Distribution Protocol (LDP) internal topology table.

Action

On Device R0, from operational mode, run the **show ldp route** command to display the internal topology table of LDP.

```
user@R0> show ldp route
```

Destination	Next-hop intf/lsp/table	Next-hop address
10.4.0.0/16	fxp0.0	10.92.31.254
10.5.0.0/16	fxp0.0	10.92.31.254
10.6.128.0/17	fxp0.0	10.92.31.254

10.9.0.0/16	fxp0.0	10.92.31.254
10.10.0.0/16	fxp0.0	10.92.31.254
10.13.4.0/23	fxp0.0	10.92.31.254
10.13.10.0/23	fxp0.0	10.92.31.254
10.82.0.0/15	fxp0.0	10.92.31.254
10.84.0.0/16	fxp0.0	10.92.31.254
10.85.12.0/22	fxp0.0	10.92.31.254
10.92.0.0/16	fxp0.0	10.92.31.254
10.92.16.0/20	fxp0.0	
10.92.20.175/32		
10.94.0.0/16	fxp0.0	10.92.31.254
10.99.0.0/16	fxp0.0	10.92.31.254
10.102.0.0/16	fxp0.0	10.92.31.254
10.150.0.0/16	fxp0.0	10.92.31.254
10.155.0.0/16	fxp0.0	10.92.31.254
10.157.64.0/19	fxp0.0	10.92.31.254
10.160.0.0/16	fxp0.0	10.92.31.254
10.204.0.0/16	fxp0.0	10.92.31.254
10.205.0.0/16	fxp0.0	10.92.31.254
10.206.0.0/16	fxp0.0	10.92.31.254
10.207.0.0/16	fxp0.0	10.92.31.254
10.209.0.0/16	fxp0.0	10.92.31.254
10.212.0.0/16	fxp0.0	10.92.31.254
10.213.0.0/16	fxp0.0	10.92.31.254
10.214.0.0/16	fxp0.0	10.92.31.254
10.215.0.0/16	fxp0.0	10.92.31.254
10.216.0.0/16	fxp0.0	10.92.31.254
10.218.13.0/24	fxp0.0	10.92.31.254
10.218.14.0/24	fxp0.0	10.92.31.254
10.218.16.0/20	fxp0.0	10.92.31.254
10.218.32.0/20	fxp0.0	10.92.31.254
10.227.0.0/16	fxp0.0	10.92.31.254
10.255.111.0/24	ge-0/0/2.0	11.11.11.2
10.255.111.4/32	ge-0/0/2.0	11.11.11.2
10.255.112.1/32	lo0.0	
10.255.112.2/32	ge-0/0/2.0	11.11.11.2
11.11.11.0/24	ge-0/0/2.0	
11.11.11.1/32		
12.12.12.0/24	ge-0/0/2.0	11.11.11.2
15.15.15.0/24	ge-0/0/1.0	
15.15.15.1/32		
22.22.22.0/24	ge-0/0/0.0	
22.22.22.1/32		
23.23.23.0/24	ge-0/0/2.0	11.11.11.2

24.24.24.0/24	ge-0/0/2.0	11.11.11.2
25.25.25.0/24	ge-0/0/2.0	11.11.11.2
128.92.17.45/32	ge-0/0/2.0	11.11.11.2
128.92.20.175/32	lo0.0	
128.92.21.186/32	ge-0/0/2.0	11.11.11.2
128.92.25.135/32	ge-0/0/2.0	11.11.11.2
128.92.27.91/32	ge-0/0/2.0	11.11.11.2
128.92.28.70/32	ge-0/0/2.0	11.11.11.2
172.16.0.0/12	fxp0.0	10.92.31.254
192.168.0.0/16	fxp0.0	10.92.31.254
192.168.102.0/23	fxp0.0	10.92.31.254
207.17.136.0/24	fxp0.0	10.92.31.254
207.17.136.192/32	fxp0.0	10.92.31.254
207.17.137.0/24	fxp0.0	10.92.31.254
224.0.0.5/32		

Meaning

The output displays the route entries in the Label Distribution Protocol (LDP) internal topology table of Device R0.

Verify Only FEC Information of LDP Route

Purpose

Display only the FEC information of LDP route.

Action

On Device R0, from operational mode, run the **show ldp route fec-only** command to display the routes in the routing table.

```
user@R0> show ldp route fec-only
```

Destination	Next-hop intf/lsp/table	Next-hop address
10.255.111.1/32	ge-0/0/2.0	11.11.11.2
10.255.111.2/32	ge-0/0/2.0	11.11.11.2
10.255.111.3/32	ge-0/0/2.0	11.11.11.2
10.255.111.4/32	ge-0/0/2.0	11.11.11.2
10.255.112.1/32	lo0.0	
10.255.112.2/32	ge-0/0/2.0	11.11.11.2

Meaning

The output displays only the FEC routes of LDP protocol available for Device R0.

Verify FEC and Shadow Routes of LDP

Purpose

Display the FEC and the shadow routes in the routing table.

Action

On Device R0, from operational mode, run the **show ldp route fec-and-route** command to display the FEC and shadow routes in the routing table.

```
user@R0> show ldp route fec-and-route
```

Destination	Next-hop intf/lsp/table	Next-hop address
10.4.0.0/16	fxp0.0	10.92.31.254
10.5.0.0/16	fxp0.0	10.92.31.254
10.6.128.0/17	fxp0.0	10.92.31.254
10.9.0.0/16	fxp0.0	10.92.31.254
10.10.0.0/16	fxp0.0	10.92.31.254
10.13.4.0/23	fxp0.0	10.92.31.254
10.13.10.0/23	fxp0.0	10.92.31.254
10.82.0.0/15	fxp0.0	10.92.31.254
10.84.0.0/16	fxp0.0	10.92.31.254
10.85.12.0/22	fxp0.0	10.92.31.254
10.92.0.0/16	fxp0.0	10.92.31.254
10.92.16.0/20	fxp0.0	
10.92.20.175/32		
10.94.0.0/16	fxp0.0	10.92.31.254
10.99.0.0/16	fxp0.0	10.92.31.254
10.102.0.0/16	fxp0.0	10.92.31.254
10.150.0.0/16	fxp0.0	10.92.31.254
10.155.0.0/16	fxp0.0	10.92.31.254
10.157.64.0/19	fxp0.0	10.92.31.254
10.160.0.0/16	fxp0.0	10.92.31.254
10.204.0.0/16	fxp0.0	10.92.31.254
10.205.0.0/16	fxp0.0	10.92.31.254
10.206.0.0/16	fxp0.0	10.92.31.254
10.207.0.0/16	fxp0.0	10.92.31.254
10.209.0.0/16	fxp0.0	10.92.31.254
10.212.0.0/16	fxp0.0	10.92.31.254
10.213.0.0/16	fxp0.0	10.92.31.254
10.214.0.0/16	fxp0.0	10.92.31.254
10.215.0.0/16	fxp0.0	10.92.31.254
10.216.0.0/16	fxp0.0	10.92.31.254
10.218.13.0/24	fxp0.0	10.92.31.254

10.218.14.0/24	fxp0.0	10.92.31.254
10.218.16.0/20	fxp0.0	10.92.31.254
10.218.32.0/20	fxp0.0	10.92.31.254
10.227.0.0/16	fxp0.0	10.92.31.254
10.255.111.0/24	ge-0/0/2.0	11.11.11.2
10.255.111.1/32	ge-0/0/2.0	11.11.11.2
10.255.111.2/32	ge-0/0/2.0	11.11.11.2
10.255.111.3/32	ge-0/0/2.0	11.11.11.2
10.255.111.4/32	ge-0/0/2.0	11.11.11.2
10.255.111.4/32	ge-0/0/2.0	11.11.11.2
10.255.112.1/32	lo0.0	
10.255.112.1/32	lo0.0	
10.255.112.2/32	ge-0/0/2.0	11.11.11.2
10.255.112.2/32	ge-0/0/2.0	11.11.11.2
11.11.11.0/24	ge-0/0/2.0	
11.11.11.1/32		
12.12.12.0/24	ge-0/0/2.0	11.11.11.2
15.15.15.0/24	ge-0/0/1.0	
15.15.15.1/32		
22.22.22.0/24	ge-0/0/0.0	
22.22.22.1/32		
23.23.23.0/24	ge-0/0/2.0	11.11.11.2
24.24.24.0/24	ge-0/0/2.0	11.11.11.2
25.25.25.0/24	ge-0/0/2.0	11.11.11.2
128.92.17.45/32	ge-0/0/2.0	11.11.11.2
128.92.20.175/32	lo0.0	
128.92.21.186/32	ge-0/0/2.0	11.11.11.2
128.92.25.135/32	ge-0/0/2.0	11.11.11.2
128.92.27.91/32	ge-0/0/2.0	11.11.11.2
128.92.28.70/32	ge-0/0/2.0	11.11.11.2
172.16.0.0/12	fxp0.0	10.92.31.254
192.168.0.0/16	fxp0.0	10.92.31.254
192.168.102.0/23	fxp0.0	10.92.31.254
207.17.136.0/24	fxp0.0	10.92.31.254
207.17.136.192/32	fxp0.0	10.92.31.254
207.17.137.0/24	fxp0.0	10.92.31.254
224.0.0.5/32		

Meaning

The output displays the FEC and the shadow routes of Device R0

RELATED DOCUMENTATION

[Longest Match Support for LDP Overview | 933](#)

[longest-match | 2667](#)

[Configuring Longest Match for LDP | 939](#)

Configuring LDP Route Preferences

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

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

```
preference preference;
```

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

LDP Graceful Restart

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

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

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

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

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

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.
- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

Configuring LDP Graceful Restart

IN THIS SECTION

- [Enabling Graceful Restart | 960](#)
- [Disabling LDP Graceful Restart or Helper Mode | 961](#)
- [Configuring Reconnect Time | 962](#)
- [Configuring Recovery Time and Maximum Recovery Time | 962](#)

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

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]

NOTE: ACX Series routers do not support [edit logical-systems *logical-system-name* routing-options] hierarchy level.

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*.

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

Disabling LDP Graceful Restart or Helper Mode

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

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

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

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

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

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

The following LDP graceful restart configurations are possible:

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

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

Configuring Reconnect Time

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

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

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

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

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

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

Configuring Recovery Time and Maximum Recovery Time

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

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

Router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from Router A.

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

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

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

Filtering Inbound LDP Label Bindings

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

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

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

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. [Table 25 on page 963](#) lists the only **from** operators that apply to LDP received-label filtering.

Table 25: from Operators That Apply to LDP Received-Label Filtering

from Operator	Description
interface	Matches on bindings received from a neighbor that is adjacent over the specified interface
neighbor	Matches on bindings received from the specified LDP router ID
next-hop	Matches on bindings received from a neighbor advertising the specified interface address
route-filter	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 User Guide*.

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 26 on page 966](#).

Table 26: to Operators for LDP Outbound-Label Filtering

to Operator	Description
interface	Matches on bindings sent to a neighbor that is adjacent over the specified interface
neighbor	Matches on bindings sent to the specified LDP router ID
next-hop	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 User Guide*.

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](#) | 2994

Control Transport Address Used for Targeted-LDP Session

IN THIS SECTION

- [Benefits of Controlling Transport Address Used for Targeted-LDP Session | 969](#)
- [Targeted-LDP Transport Address Overview | 969](#)
- [Transport Address Preference | 970](#)
- [Troubleshooting Transport Address Configuration | 970](#)

To establish a TCP session between two devices, each device must learn the other device's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session operates. Earlier, this transport address could only be the router-ID or an interface address. With the LDP transport-address feature, you can explicitly configure any IP address as the transport address for targeted LDP neighbors for Layer 2 circuit, MPLS, and VPLS adjacencies. This enables you to control the targeted-LDP sessions using transport-address configuration.

Benefits of Controlling Transport Address Used for Targeted-LDP Session

Configuring transport address for establishing targeted-LDP sessions has the following benefits:

- **Flexible interface configurations**—Provides the flexibility of configuring multiple IP addresses for one loopback interface without interrupting the creation of LDP session between the targeted-LDP neighbors.
- **Ease of operation**—Transport address configured at the interface-level, allows you to use more than one protocol in the IGP backbone for LDP. This enables smooth and easy operations.

Targeted-LDP Transport Address Overview

Prior to Junos OS Release 19.1R1, LDP provided support only for router-ID or the interface address as the transport address on any LDP interface. The adjacencies formed on that interface used one of the IP addresses assigned to the interface or the router-ID. In case of targeted adjacency, the interface is the loopback interface. When multiple loopback addresses were configured on the device, the transport address could not be derived for the interface, and as a result, the LDP session could not be established.

Starting in Junos OS Release 19.1R1, in addition to the default IP addresses used for transport address of targeted-LDP sessions, you can configure any other IP address as the transport address under the **session**, **session-group**, and **interface** configuration statements. The transport address configuration is applicable for configured neighbors only including Layer 2 circuits, MPLS, and VPLS adjacencies. This configuration does not apply to discovered adjacencies (targeted or not).

Transport Address Preference

You can configure transport address for targeted-LDP sessions at the session, session-group, and interface level.

After the transport address is configured, the targeted-LDP session is established based on the transport address preference of LDP.

The order of preference of transport address for targeted neighbor (configured through Layer 2 circuit, MPLS, VPLS, and LDP configuration) is as follows:

1. Under **[edit protocols ldp session]** hierarchy.
2. Under **[edit protocols ldp session-group]** hierarchy.
3. Under **[edit protocols ldp interface lo0]** hierarchy.
4. Under **[edit protocols ldp]** hierarchy.
5. Default address.

The order of preference of transport address for the discovered neighbors is as follows:

1. Under **[edit protocols ldp interface]** hierarchy.
2. Under **[edit protocols ldp]** hierarchy.
3. Default address.

The order of preference of transport address for auto-targeted neighbors where LDP is configured to accept hello packets is as follows:

1. Under **[edit protocols ldp interface lo0]** hierarchy.
2. Under **[edit protocols ldp]** hierarchy.
3. Default address.

Troubleshooting Transport Address Configuration

You can use the following show command outputs to troubleshoot targeted-LDP sessions:

- **show ldp session**
- **show ldp neighbor**

The **detail** level of output of the **show ldp neighbor** command displays the transport address sent in the hello messages to the targeted neighbor. If this address is not reachable from the neighbor, the LDP session does not come up.

- **show configuration protocols ldp**

You can also enable LDP traceoptions for further troubleshooting.

- If the configuration is changed from using a transport address that is invalid (non reachable) to transport address that is valid, the following traces can be observed:

```
May 29 10:47:11.569722 Incoming connect from 10.55.1.4
May 29 10:47:11.570064 Connection 10.55.1.4 state Closed -> Open
May 29 10:47:11.570727 Session 10.55.1.4 state Nonexistent -> Initialized
May 29 10:47:11.570768 Session 10.55.1.4 state Initialized -> OpenRec
May 29 10:47:11.570799 LDP: Session param Max PDU length 4096 from 10.55.1.4,
negotiated 4096
May 29 10:47:11.570823 Session 10.55.1.4 GR state Nonexistent -> Operational
May 29 10:47:11.669295 Session 10.55.1.4 state OpenRec -> Operational
May 29 10:47:11.669387 RPD_LDP_SESSIONUP: LDP session 10.55.1.4 is up
```

- If the configuration is changed from using a transport address that is valid to transport address that is invalid (non reachable), the following traces can be observed:

```
May 29 10:42:36.317942 Session 10.55.1.4 GR state Operational -> Nonexistent
May 29 10:42:36.318171 Session 10.55.1.4 state Operational -> Closing
May 29 10:42:36.318208 LDP session 10.55.1.4 is down, reason: received notification
from peer
May 29 10:42:36.318236 RPD_LDP_SESSIONDOWN: LDP session 10.55.1.4 is down, reason:
received notification from peer
May 29 10:42:36.320081 Connection 10.55.1.4 state Open -> Closed
May 29 10:42:36.322411 Session 10.55.1.4 state Closing -> Nonexistent
```

In case of faulty configuration, perform the following troubleshooting tasks:

- Check the **address family**. The transport address that is configured under the **session** statement must belong to the same address family as the neighbor or session.
- The address that is configured as the transport address under a **neighbor** or **session** statement must be local to the router for the targeted hello messages to start. You can check if the address is configured. If the address is not configured under any interface, the configuration is rejected.

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*.

NOTE: ACX Series routers do not support **[edit logical-systems]** hierarchy level.

Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

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


Configuring FEC Deaggregation

When an LDP egress router advertises multiple prefixes, the prefixes are bound to a single label and aggregated into a single forwarding equivalence class (FEC). By default, LDP maintains this aggregation as the advertisement traverses the network.

Normally, because an LSP is not split across multiple next hops and the prefixes are bound into a single LSP, load-balancing across equal-cost paths does not occur. You can, however, load-balance across equal-cost paths if you configure a load-balancing policy and deaggregate the FECs.

Deaggregating the FECs causes each prefix to be bound to a separate label and become a separate LSP.

To configure deaggregated FECs, include the **deaggregate** statement:

```
deaggregate;
```

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

For all LDP sessions, you can configure deaggregated FECs only globally.

Deaggregating a FEC allows the resulting multiple LSPs to be distributed across multiple equal-cost paths and distributes LSPs across the multiple next hops on the egress segments but installs only one next hop per LSP.

To aggregate FECs, include the **no-deaggregate** statement:

```
no-deaggregate;
```

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

For all LDP sessions, you can configure aggregated FECs only globally.

RELATED DOCUMENTATION

[Configuring Load Balancing Across RSVP LSPs | 848](#)

Configuring Protocol-Independent Load Balancing in Layer 3 VPNs

Configuring VPLS Load Balancing

Example: Load Balancing BGP Traffic

Configuring Policers for LDP FECs

You can configure the Junos OS to track and police traffic for LDP FECs. LDP FEC policers can be used to do any of the following:

- Track or police the ingress traffic for an LDP FEC.
- Track or police the transit traffic for an LDP FEC.
- Track or police LDP FEC traffic originating from a specific forwarding class.
- Track or police LDP FEC traffic originating from a specific virtual routing and forwarding (VRF) site.
- Discard false traffic bound for a specific LDP FEC.

To police traffic for an LDP FEC, you must first configure a filter. Specifically, you need to configure either the **interface** statement or the **interface-set** statement at the **[edit firewall family protocol-family filter filter-name term term-name from]** hierarchy level. The **interface** statement allows you to match the filter to a single interface. The **interface-set** statement allows you to match the filter to multiple interfaces.

For more information on how to configure the **interface** statement, the **interface-set** statement, and policers for LDP FECs, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

Once you have configured the filters, you need to include them in the **policing** statement configuration for LDP. To configure policers for LDP FECs, include the **policing** statement:

```
policing {  
  fec fec-address {  
    ingress-traffic filter-name;  
    transit-traffic filter-name;  
  }  
}
```

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

The **policing** statement includes the following options:

- **fec**—Specify the FEC address for the LDP FEC you want to police.
- **ingress-filter**—Specify the name of the ingress traffic filter.
- **transit-traffic**—Specify the name of the transit traffic filter.

Configuring LDP IPv4 FEC Filtering

By default, when a targeted LDP session is established, the Junos OS always exchanges both the IPv4 forwarding equivalence classes (FECs) and the Layer 2 circuit FECs over the targeted LDP session. For an LDP session to an indirectly connected neighbor, you might only want to export Layer 2 circuit FECs to the neighbor if the session was specifically configured to support Layer 2 circuits or VPLS.

In a mixed vendor network where all non-BGP prefixes are advertised into LDP, the LDP database can become large. For this type of environment, it can be useful to prevent the advertisement of IPv4 FECs over LDP sessions formed because of Layer 2 circuit or LDP VPLS configuration. Similarly, it can be useful to filter any IPv4 FECs received in this sort of environment.

If all the LDP neighbors associated with an LDP session are Layer 2 only, you can configure the Junos OS to advertise only Layer 2 circuit FECs by configuring the **I2-smart-policy** statement. This feature also automatically filters out the IPv4 FECs received on this session. Configuring an explicit export or import policy that is activated for **I2-smart-policy** disables this feature in the corresponding direction.

If one of the LDP session's neighbors is formed because of a discovered adjacency or if the adjacency is formed because of an LDP tunneling configuration on one or more RSVP LSPs, the IPv4 FECs are advertised and received using the default behavior.

To prevent LDP from exporting IPv4 FECs over LDP sessions with Layer 2 neighbors only and to filter out IPv4 FECs received over such sessions, include the **I2-smart-policy** statement:

```
I2-smart-policy;
```

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

Configuring BFD for LDP LSPs

You can configure Bidirectional Forwarding Detection (BFD) for LDP LSPs. The BFD protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the router stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of static routes, providing faster detection.

An error is logged whenever a BFD session for a path fails. The following shows how BFD for LDP LSP log messages might appear:

```
RPD_LDP_BFD_UP: LDP BFD session for FEC 10.255.16.14/32 is up
RPD_LDP_BFD_DOWN: LDP BFD session for FEC 10.255.16.14/32 is down
```

You can also configure BFD for RSVP LSPs, as described in [“Configuring BFD for RSVP-Signaled LSPs” on page 110](#).

The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

To enable BFD for LDP LSPs, include the **oam** and **bfd-liveness-detection** statements:

```
oam {
  bfd-liveness-detection {
    detection-time threshold milliseconds;
    ecmp;
    failure-action {
      remove-nexthop;
      remove-route;
    }
    holddown-interval seconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (0 | 1 | automatic);
  }
  fec fec-address {
    bfd-liveness-detection {
```



```

    detection-time threshold milliseconds;
    ecmp;
    failure-action {
        remove-nexthop;
        remove-route;
    }
    holddown-interval milliseconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (0 | 1 | automatic);
}
no-bfd-liveness-detection;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
}
lsp-ping-interval seconds;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}

```



```
}
```

You can enable BFD for the LDP LSPs associated with a specific forwarding equivalence class (FEC) by configuring the FEC address using the **fec** option at the **[edit protocols ldp]** hierarchy level. Alternatively, you can configure an Operation Administration and Management (OAM) ingress policy to enable BFD on a range of FEC addresses. For more information, see [“Configuring OAM Ingress Policies for LDP” on page 1238](#).

You cannot enable BFD LDP LSPs unless their equivalent FEC addresses are explicitly configured or OAM is enabled on the FECs using an OAM ingress policy. If BFD is not enabled for any FEC addresses, the BFD session will not come up.

You can configure the **oam** statement at the following hierarchy levels:

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

NOTE: ACX Series routers do not support **[edit logical-systems]** hierarchy level.

The **oam** statement includes the following options:

- **fec**—Specify the FEC address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.
- **lsp-ping-interval**—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command. For more information, see the [CLI Explorer](#).

The **bfd-liveness-detection** statement includes the following options:

- **ecmp**—Cause LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the **ecmp** option, you must also configure the **periodic-traceroute** statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the **periodic-traceroute** statement at the global hierarchy level (**[edit protocols ldp oam]**) while only configuring the **ecmp** option for a specific FEC (**[edit protocols ldp oam fec address bfd-liveness-detection]**).
- **holddown-interval**—Specify the duration the BFD session should remain up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.
- **minimum-interval**—Specify the minimum transmit and receive interval. If you configure the **minimum-interval** option, you do not need to configure the **minimum-receive-interval** option or the **minimum-transmit-interval** option.
- **minimum-receive-interval**—Specify the minimum receive interval. The range is from 1 through 255,000 milliseconds.

- **minimum-transmit-interval**—Specify the minimum transmit interval. The range is from 1 through 255,000 milliseconds.
- **multiplier**—Specify the detection time multiplier. The range is from 1 through 255.
- **version**—Specify the BFD version. The options are BFD version 0 or BFD version 1. By default, the Junos OS software attempts to automatically determine the BFD version.

Configuring ECMP-Aware BFD for LDP LSPs

When you configure BFD for a FEC, a BFD session is established for only one active local next-hop for the router. However, you can configure multiple BFD sessions, one for each FEC associated with a specific equal-cost multipath (ECMP) path. For this to function properly, you also need to configure LDP LSP periodic traceroute. (See [“Configuring LDP LSP Traceroute” on page 1123](#).) LDP LSP traceroute is used to discover ECMP paths. A BFD session is initiated for each ECMP path discovered. Whenever a BFD session for one of the ECMP paths fails, an error is logged.

LDP LSP traceroute is run periodically to check the integrity of the ECMP paths. The following might occur when a problem is discovered:

- If the latest LDP LSP traceroute for a FEC differs from the previous traceroute, the BFD sessions associated with that FEC (the BFD sessions for address ranges that have changed from previous run) are brought down and new BFD sessions are initiated for the destination addresses in the altered ranges.
- If the LDP LSP traceroute returns an error (for example, a timeout), all the BFD sessions associated with that FEC are torn down.

To configure LDP to establish BFD sessions for all ECMP paths configured for the specified FEC, include the **ecmp** statement.

```
ecmp;
```

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

Along with the **ecmp** statement, you must also include the **periodic-traceroute** statement, either in the global LDP OAM configuration (at the **[edit protocols ldp oam]** or **[edit logical-systems logical-system-name protocols ldp oam]** hierarchy level) or in the configuration for the specified FEC (at the **[edit protocols ldp oam fec address]** or **[edit logical-systems logical-system-name protocols ldp oam fec address]** hierarchy level). Otherwise, the commit operation fails.

NOTE: ACX Series routers do not support **[edit logical-systems]** hierarchy level.

Configuring a Failure Action for the BFD Session on an LDP LSP

You can configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.

You can configure one of the following failure action options for the **failure-action** statement in the event of a BFD session failure on the LDP LSP:

- **remove-nexthop**—Removes the route corresponding to the next hop of the LSP's route at the ingress node when a BFD session failure event is detected.
- **remove-route**—Removes the route corresponding to the LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

To configure a failure action in the event of a BFD session failure on an LDP LSP, include either the **remove-nexthop** option or the **remove-route** option for the **failure-action** statement:

```
failure-action {
  remove-nexthop;
  remove-route;
}
```

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

Configuring the Holddown Interval for the BFD Session

You can specify the duration the BFD session should be up before adding a route or next hop by configuring the **holddown-interval** statement at either the **[edit protocols ldp oam bfd-liveness-detection]** hierarchy level or at the **[edit protocols ldp oam fec address bfd-liveness-detection]** hierarchy level. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

```
holddown-interval seconds;
```

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

Configuring LDP Link Protection

You can configure Label Distribution Protocol (LDP) link protection for both unicast and multicast LDP label-switched paths (LSPs) to provide resiliency during link or node failure.

Before you begin:

1. Configure the device interfaces.
2. Configure the router ID and autonomous system number for the device.
3. Configure the following protocols:
 - a. RSVP
 - b. MPLS with traffic engineering capability.
 - c. OSPF with traffic engineering capability.

NOTE: For multicast LDP link protection with loop-free alternative (LFA), enable link protection.

```
[edit protocols]
user@R0# set ospf area 0 interface all link-protection
```

To configure LDP link protection:

1. Enable point-to-multipoint LDP LSPs.

```
[edit protocols]
user@R0# set ldp p2mp
```

2. Enable LDP on all the interfaces of Router R0 (excluding the management interface) and configure link protection with dynamic RSVP bypass LSP.

```
[edit protocols]
user@R0# set ldp interface all link-protection dynamic-rsvp-lsp
user@R0# set ldp interface fxp0.0 disable
```

3. Verify and commit the configuration.

For example:


```
[edit protocols]
user@R0# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  traffic-engineering;
  interface all;
  interface fxp0.0 {
    disable;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all {
      metric 1;
    }
    interface fxp0.0 {
      disable;
    }
  }
}
ldp {
  interface all {
    link-protection {
      dynamic-rsvp-lsp;
    }
  }
  interface fxp0.0 {
    disable;
  }
  p2mp;
}
```

```
[edit]
user@R0# commit
commit complete
```


RELATED DOCUMENTATION

| [Example: Configuring LDP Link Protection](#) | 983

Example: Configuring LDP Link Protection

IN THIS SECTION

- [LDP Link Protection Overview](#) | 983
- [Example: Configuring LDP Link Protection](#) | 1002

LDP Link Protection Overview

IN THIS SECTION

- [Introduction to LDP](#) | 983
- [Junos OS LDP Protocol Implementation](#) | 984
- [Understanding Multipoint Extensions to LDP](#) | 984
- [Using Multipoint Extensions to LDP on Targeted LDP Sessions](#) | 984
- [Current Limitations of LDP Link Protection](#) | 985
- [Using RSVP LSP as a Solution](#) | 986
- [Understanding Multicast LDP Link Protection](#) | 988
- [Different Modes for Providing LDP Link Protection](#) | 989
- [Label Operation for LDP Link Protection](#) | 991
- [Sample Multicast LDP Link Protection Configuration](#) | 998
- [Make-Before-Break](#) | 1000
- [Caveats and Limitations](#) | 1001

Introduction to LDP

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

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

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

Junos OS LDP Protocol Implementation

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

Understanding Multipoint Extensions to LDP

An LDP defines mechanisms for setting up point-to-point, multipoint-to-point, point-to-multipoint, and multipoint-to-multipoint LSPs in the network. The point-to-multipoint and multipoint-to-multipoint LSPs are collectively referred to as multipoint LSPs, where traffic flows from a single source to multiple destinations, and from multiple sources to multiple destinations, respectively. The destination or egress routers are called leaf nodes, and traffic from the source traverses one or more transit nodes before reaching the leaf nodes.

NOTE: Junos OS does not provide support for multipoint-to-multipoint LSPs.

By taking advantage of the MPLS packet replication capability of the network, multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.


Using Multipoint Extensions to LDP on Targeted LDP Sessions

The specification for the multipoint extensions to LDP requires that the two endpoints of an LDP session are directly connected by a Layer 2 medium, or are considered to be neighbors by the network's IGP. This is referred to as an LDP link session. When the two endpoints of an LDP session are not directly connected, the session is referred to as a targeted LDP session.

Past Junos OS implementations support multicast LDP for link sessions only. With the introduction of the LDP link protection feature, the multicast LDP capabilities are extended to targeted LDP sessions.

[Figure 74 on page 985](#) shows a sample topology.

Figure 74: Multicast LDP Support for Targeted LDP Session



Routers R7 and R8 are the upstream (LSR-U) and downstream (LSR-D) label-switched routers (LSRs), respectively, and deploy multicast LDP. The core router, Router R5, has RSVP-TE enabled.

When LSR-D is setting up the point-to-multipoint LSP with root and LSP ID attributes, it determines the upstream LSR-U as a next-hop on the best path to the root (currently, this next-hop is assumed to be an IGP next hop).

With the multicast LDP support on targeted LDP sessions, you can determine if there is an LSP next hop to LSR-U which is on LSR-D's path to root, where LSR-D and LSR-U are not directly connected neighbors, but targeted LDP peers. The point-to-multipoint label advertised on the targeted LDP session between LSR-D and LSR-U is not used unless there is an LSP between LSR-D and LSR-U. Therefore, a corresponding LSP in the reverse direction from LSR-U to LSR-D is required.

Data is transmitted on the point-to-multipoint LSP using unicast replication of packets, where LSR-U sends one copy to each downstream LSR of the point-to-multipoint LSP.

The data transmission is implemented in the following ways:

1. The point-to-multipoint capabilities on the targeted LDP session are negotiated.
2. The algorithm to select the upstream LSR is changed, where if IGP next hops are unavailable, or in other words, there is no LDP link session between LSR-D and LSR-U, an RSVP LSP is used as the next hop to reach LSR-U.
3. The incoming labels received over the targeted LDP sessions are installed as a branch next hop for this point-to-multipoint FEC route with the LDP label as the inner label and the RSVP label as the outer label.

Current Limitations of LDP Link Protection

When there is a link or node failure in an LDP network deployment, fast traffic recovery should be provided to recover impacted traffic flows for mission-critical services. In the case of multipoint LSPs, when one of the links of the point-to-multipoint tree fails, the subtrees might get detached until the IGP reconverges

and the multipoint LSP is established using the best path from the downstream router to the new upstream router.


In fast reroute using local repair for LDP traffic, a backup path (repair path) is pre-installed in the Packet Forwarding Engine. When the primary path fails, traffic is rapidly moved to the backup path without having to wait for the routing protocols to converge. Loop-free alternate (LFA) is one of the methods used to provide IP fast reroute capability in the core and service provider networks.

Without LFA, when a link or a router fails or is returned to service, the distributed routing algorithms compute the new routes based on the changes in the network. The time during which the new routes are computed is referred to as routing transition. Until the routing transition is completed, the network connectivity is interrupted because the routers adjacent to a failure continue to forward the data packets through the failed component until an alternative path is identified.

However, LFA does not provide full coverage in all network deployments because of the IGP metrics. As a result, this is a limitation to the current LDP link protection schemes.

Figure 75 on page 986 illustrates a sample network with incomplete LFA coverage, where traffic flows from the source router (S) to the destination router (D) through Router R1. Assuming that each link in the network has the same metric, if the link between the Router S and Router R1 fails, Router R4 is not an LFA that protects the S-R1 link, so traffic resiliency is lost. Thus, full coverage is not achieved by using plain LFA. In typical networks, there is always some percentage of LFA coverage gap with plain LFA.

Figure 75: Incomplete Coverage Problem with LFA



Using RSVP LSP as a Solution

IN THIS SECTION

- Manually Configured RSVP LSPs | 987
- Dynamically Configured RSVP LSPs | 987


The key to protect the traffic flowing through LDP LSPs is to have an explicit tunnel to re-route the traffic in the event of a link or node failure. The explicit path has to terminate on the next downstream router, and the traffic needs to be accepted on the explicit path, where the RPF check should pass.

RSVP LSPs help overcome the current limitations of loop-free alternate (LFA) for both point-to-point and point-to-multipoint LDP LSPs by extending the LFA coverage in the following methods:

Manually Configured RSVP LSPs

Considering the example used in [Figure 75 on page 986](#), when the S-R1 link fails, and Router R4 is not an LFA for that particular link, a manually created RSVP LSP is used as a patch to provide complete LFA coverage. The RSVP LSP is pre-sigaled and pre-installed in the Packet Forwarding Engine of Router S, so that it can be used as soon as Router S detects that the link has failed.

Figure 76: Manually Configured RSVP LSP Coverage



In this case, an RSVP LSP is created between Routers S, R4, and R3 as illustrated in [Figure 76 on page 987](#). A targeted LDP session is created between Router S and Router R3, as a result of which, when the S-R1 link fails, traffic reaches Router R3. Router R3 forwards the traffic to Router R2, as it is the shortest path to reach the destination, Router D.


Dynamically Configured RSVP LSPs

In this method, the RSVP LSPs are created automatically and pre-installed in the system so that they can be used immediately when there is a link failure. Here, the egress is the node on the other side of the link being protected, thereby improving the LFA coverage.

Considering the example used in [Figure 75 on page 986](#), in order to protect traffic against the potential failure of the S-R1 link, because Router R4 is not an LFA for that particular link, an RSVP bypass LSP is automatically created to Router R1, which is the node on the far side of the protected link as illustrated in [Figure 77 on page 988](#). From Router R1, traffic is forwarded to its original destination, Router D.

The RSVP LSP is pre-sigaled and pre-installed in the Packet Forwarding Engine of Router S so that it can be used as soon as Router S detects that the link has failed.

Figure 77: Dynamically Configured RSVP LSP Coverage



An alternative mode of operation is not to use LFA at all, and to always have the RSVP LSP created to cover all link failures.

To enable dynamic RSVP LSPs, include the **dynamic-rsvp-lsp** statement at the **[edit protocols ldp interface interface-name link-protection]** hierarchy level, in addition to enabling the RSVP protocol on the appropriate interfaces.

Some of the benefits of enabling dynamic RSVP LSPs include:

- Ease of configuration.
- 100 percent coverage against link failure as long as there is an alternate path to the far end of the link being protected.
- Setting up and tearing down of the RSVP bypass LSP is automatic.
- RSVP LSP only used for link protection and not for forwarding traffic while the link being protected is up.
- Reduces the total number of RSVP LSPs required on the system.

Understanding Multicast LDP Link Protection

A point-to-multipoint LDP label-switched path (LSP) is an LDP-sigaled LSP that is point-to-multipoint, and is referred to as multicast LDP.

A multicast LDP LSP can be used to send traffic from a single root or ingress node to a number of leaf or egress nodes traversing one or more transit nodes. Multicast LDP link protection enables fast reroute of traffic carried over point-to-multipoint LDP LSPs in case of a link failure. When one of the links of the point-to-multipoint tree fails, the subtrees might get detached until the IGP reconverges and the multipoint LSP is established using the best path from the downstream router to the new upstream router.

To protect the traffic flowing through the multicast LDP LSP, you can configure an explicit tunnel to re-route the traffic in the event of link failure. The explicit path has to terminate on the next downstream router. The reverse path forwarding for the traffic should be successful.

Multicast LDP link protection introduces the following features and functionality:

- Use of dynamic RSVP LSP as bypass tunnels

The RSVP LSP's Explicit Route Object (ERO) is calculated using Constrained Shortest Path First (CSPF) with the constraint as the link to avoid. The LSP is signaled and torn down dynamically whenever link protection is necessary.

- Make-before-break

The make-before-break feature ensures that there is minimum packet loss when attempting to signal a new LSP path before tearing down the old LSP path for the multicast LDP LSP.

- Targeted LDP session

A targeted adjacency to the downstream label-switching router (LSR) is created for two reasons:

- To keep the session up after link failure.
- To use the point-to-multipoint label received from the session to send traffic to the downstream LSR on the RSVP LSP bypass tunnel.

When the downstream LSR sets up the multicast LDP LSP with the root node and LSP ID, it uses that upstream LSR, which is on the best path toward the root.

NOTE: Multicast LDP link protection is not required when there are multiple link adjacencies (parallel links) to the downstream LSR.

Different Modes for Providing LDP Link Protection

Following are three different modes of operation available for unicast and multicast LDP link protection:

- **Case A: LFA only**

Under this mode of operation, multicast LDP link protection is provided using an existing viable loop-free alternate (LFA). In the absence of a viable LFA, link protection is not provided for the multicast LDP LSP.

- **Case B: LFA and Dynamic RSVP LSP**


Under this mode of operation, multicast LDP link protection is provided using an existing viable LFA. In the absence of a viable LFA, an RSVP bypass LSP is created automatically to provide link protection for the multicast LDP LSP.

- **Case C: Dynamic RSVP LSP only**

Under this mode of operation, LFA is not used for link protection. Multicast LDP link protection is provided by using automatically created RSVP bypass LSP.

Figure 78 on page 990 is a sample topology illustrating the different modes of operation for multicast LDP link protection. Router R5 is the root connecting to two leaf nodes, Routers R3 and R4. Router R0 and Router R1 are the upstream and downstream label-switched routers (LSRs), respectively. A multicast LDP LSP runs among the root and leaf nodes.

Figure 78: Multicast LDP Link Protection Sample Topology



Considering that Router R0 needs to protect the multicast LDP LSP in the case that the R0-R1 link fails, the different modes of link protection operate in the following manner:

- **Case A: LFA only**

Router R0 checks if a viable LFA path exists that can avoid the R0-R1 link to reach Router R1. Based on the metrics, Router R2 is a valid LFA path for the R0-R1 link and is used to forward unicast LDP traffic. If multiple multicast LDP LSPs use the R0-R1 link, the same LFA (Router R2) is used for multicast LDP link protection.

When the R0-R1 link fails, the multicast LDP LSP traffic is moved onto the LFA path by Router R0, and the unicast LDP label to reach Router R1 (L100) is pushed on top of the multicast LDP label (L21).

- **Case B: LFA and Dynamic RSVP LSP**

Router R0 checks if a viable LFA path exists that can avoid the R0-R1 link to reach Router R1. Based on the metrics, Router R2 is a valid LFA path for the R0-R1 link and is used to forward unicast LDP traffic. If multiple multicast LDP LSPs use the R0-R1 link, the same LFA (Router R2) is used for multicast LDP

link protection. When the R0-R1 link fails, the multicast LDP LSP traffic is moved onto the LFA path by Router R0.

However, if the metric on the R2-R1 link was 50 instead of 10, Router 2 is not a valid LFA for the R0-R1 link. In this case, an RSVP LSP is automatically created to protect the multicast LDP traffic traveling between Routers R0 and R1.

- **Case C: Dynamic RSVP LSP only**

An RSVP LSP is signaled automatically from Router R0 to Router R1 through Router R2, avoiding interface ge-1/1/0. If multiple multicast LDP LSPs use the R0-R1 link, the same RSVP LSP is used for multicast LDP link protection.

When the R0-R1 link fails, the multicast LDP LSP traffic is moved onto the RSVP LSP by Router R0, and the RSVP label to reach Router R1 (L100) is pushed on top of the multicast LDP label (L21).


Label Operation for LDP Link Protection

IN THIS SECTION

- Case A: LFA Only | [992](#)
- Case B: LFA and Dynamic RSVP LSP | [996](#)
- Case C: Dynamic RSVP LSP Only | [998](#)

Using the same network topology as in Figure 5, [Figure 79 on page 992](#) illustrates the label operation for unicast and multicast LDP link protection.

Figure 79: LDP Label Operation Sample Topology



Router R5 is the root connecting to two leaf nodes, Routers R3 and R4. Router R0 and Router R1 are the upstream and downstream label-switched routers (LSRs), respectively. A multicast LDP LSP runs among the root and leaf nodes. An unicast LDP path connects Router R1 to Router R5.

The label operation is performed differently under the following modes of LDP link protection:

Case A: LFA Only

Using the **show route detail** command output on Router R0, the unicast LDP traffic and multicast LDP traffic can be derived.

```
user@R0> show route detail
```

```
299840 (1 entry, 1 announced)
  *LDP      Preference: 9
            Next hop type: Router
            Address: 0x93bc22c
            Next-hop reference count: 1
            Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1, selected
            Label operation: Swap 299824
            Session Id: 0x1
            Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0xf000
            Label operation: Swap 299808
            Session Id: 0x3
            State: <Active Int>
            Age: 3:16      Metric: 1
            Validation State: unverified
            Task: LDP
```



```

Announcement bits (1): 0-KRT
AS path: I
Prefixes bound to route: 192.168.0.4/32

299856 (1 entry, 1 announced)
  *LDP    Preference: 9
          Next hop type: Flood
          Address: 0x9340e04
          Next-hop reference count: 3
          Next hop type: Router, Next hop index: 262143
          Address: 0x93bc3dc
          Next-hop reference count: 2
          Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1
          Label operation: Swap 299888
          Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0xf000
          Label operation: Swap 299888, Push 299776(top)
          Label TTL action: prop-ttl, prop-ttl(top)
          State: <Active Int AckRequest>
          Age: 3:16          Metric: 1
          Validation State: unverified
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
          FECs bound to route: P2MP root-addr 192.168.0.5, lsp-id 99

```

Label 299840 is traffic arriving at Router R0 that corresponds to unicast LDP traffic to Router R1. Label 299856 is traffic arriving at Router 0 that corresponds to multicast LDP traffic from the root node R5 to the leaf egress nodes, R3 and R4.

The main path for both unicast and multicast LDP LSPs is through interface ge-0/0/1 (the link to Router R1), and the LFA path is through interface ge-0/0/2 (the link to Router R2). The LFA path is not used unless the ge-0/0/1 interface goes down.


In the label operation for Case A, the LFA-only mode of operation is different for unicast and multicast LDP traffic:

- Unicast label operation

For unicast LDP traffic, the FECs and associated labels are advertised on all the links in the network on which LDP is enabled. This means that in order to provide LFA action for the unicast LDP traffic to Router R4, instead of swapping the incoming label for label 299824 advertised by Router R1 for FEC R4, Router R0 simply swaps the incoming label for label 299808 advertised by Router R2 for FEC R4. This is the standard Junos OS LFA operation for unicast LDP traffic.

Figure 80 on page 994 illustrates the label operation for unicast traffic when the R0-R1 link fails. The grey boxes show the label operation for unicast LDP traffic under normal condition, and the dotted boxes show the label operation for unicast LDP traffic when the R0-R1 link fails.

Figure 80: Unicast LDP Label Operation



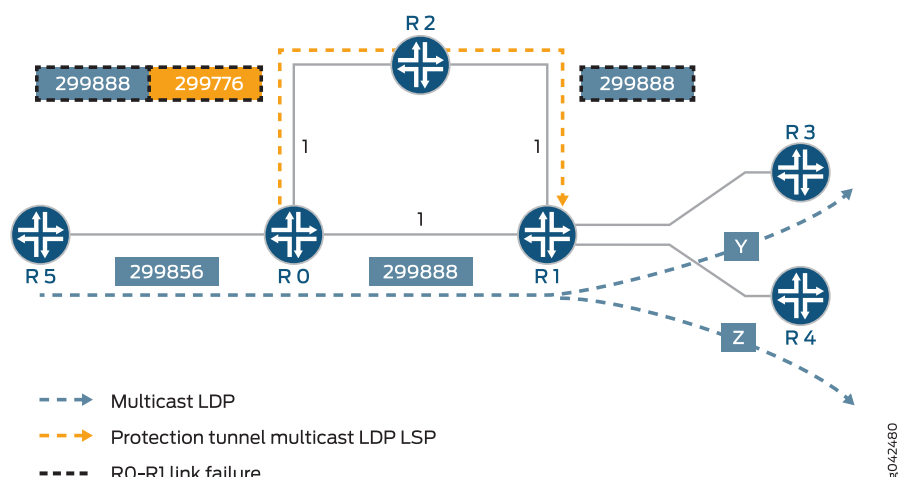
- Multicast label operation

The label operation for multicast LDP traffic differs from the unicast LDP label operation, because multipoint LSP labels are only advertised along the best path from the leaf node to the ingress node. As a result, Router R2 has no knowledge of the multicast LDP. To overcome this, the multicast LDP LSP traffic is simply tunneled inside the unicast LDP LSP path through Router R2 that terminates at Router R1.

In order to achieve this, Router R0 first swaps the incoming multicast LDP LSP label 299856 to label 299888 advertised by Router R1. Label 299776 is then pushed on top, which is the LDP label advertised by Router R2 for FEC R1. When the packet arrives at Router R2, the top label is popped out due to penultimate hop-popping. This means that the packet arrives at Router R1 with the multicast LDP label 299888 that Router R1 had originally advertised to Router R0.

Figure 81 on page 995 illustrates the label operation for multicast LDP traffic when the R0-R1 link fails. The blue boxes show the label operation for multicast LDP traffic under normal condition, and the dotted boxes show the label operation for multicast LDP traffic when the R0-R1 link fails.

Figure 81: Multicast LDP Label Operation



When the metric on the R2-R1 link is set to 1000 instead of 1, Router R2 is not a valid LFA for Router R0. In this case, if Router R2 receives a packet destined for Router R1, R3, or R4 before its IGP has converged, the packet is sent back to Router R0, resulting in looping packets.

Because Router R0 has no viable LFA, no backup paths are installed in the Packet Forwarding Engine. If the R0-R1 link fails, traffic flow is interrupted until the IGP and LDP converge and new entries are installed on the affected routers.

The **show route detail** command displays the state when no LFA is available for link protection.

```
user@host> show route detail
```

```
299840 (1 entry, 1 announced)
  *LDP      Preference: 9
            Next hop type: Router, Next hop index: 578
            Address: 0x9340d20
            Next-hop reference count: 2
            Next hop: 11.0.0.6 via ge-0/0/1.0, selected
            Label operation: Swap 299824
            Session Id: 0x1
            State: <Active Int>
            Age: 5:38      Metric: 1
            Validation State: unverified
            Task: LDP
            Announcement bits (1): 0-KRT
            AS path: I
            Prefixes bound to route: 192.168.0.4/32
```



```

299856 (1 entry, 1 announced)
  *LDP      Preference: 9
            Next hop type: Flood
            Address: 0x9340e04
            Next-hop reference count: 3
            Next hop type: Router, Next hop index: 579
            Address: 0x93407c8
            Next-hop reference count: 2
            Next hop: 11.0.0.6 via ge-0/0/1.0
            Label operation: Swap 299888
            State: <Active Int AckRequest>
            Age: 5:38      Metric: 1
            Validation State: unverified
            Task: LDP
            Announcement bits (1): 0-KRT
            AS path: I
            FECs bound to route: P2MP root-addr 192.168.0.5, lsp-id 99

```

Case B: LFA and Dynamic RSVP LSP

In this mode of operation, if there is a viable LFA neighbor, the label operation behavior is similar to that of Case A, LFA only mode. However, if there is no viable LFA neighbor, an RSVP bypass tunnel is automatically created.

If the metric on the link R2-R1 is set to 1000 instead of 1, Router R2 is not an LFA for Router R0. On learning that there are no LFA paths to protect the R0-R1 link failure, an RSVP bypass tunnel is automatically created with Router R1 as the egress node and follows a path that avoids the R0-R1 link (for instance, R0-R2-R1).

If the R0-R1 link fails, the unicast LDP and multicast LDP traffic is tunneled through the RSVP bypass tunnel. The RSVP bypass tunnel is not used for normal forwarding and is used only to provide link protection to LDP traffic in the case of R0-R1 link failure.

Using the **show route detail** command, the unicast and multicast LDP traffic can be derived.

```
user@host> show route detail
```

```

299840 (1 entry, 1 announced)
  *LDP      Preference: 9
            Next hop type: Router
            Address: 0x940c3dc
            Next-hop reference count: 1
            Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1, selected
            Label operation: Swap 299824

```



```

Session Id: 0x1
Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0x8001
Label-switched-path ge-0/0/1.0:BypassLSP->192.168.0.1
Label operation: Swap 299824, Push 299872(top)
Label TTL action: prop-ttl, prop-ttl(top)
Session Id: 0x3
State: <Active Int NhAckRequest>
Age: 19          Metric: 1
Validation State: unverified
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
Prefixes bound to route: 192.168.0.4/32

299856 (1 entry, 1 announced)
  *LDP    Preference: 9
          Next hop type: Flood
          Address: 0x9340e04
          Next-hop reference count: 3
          Next hop type: Router, Next hop index: 262143
          Address: 0x940c154
          Next-hop reference count: 2
          Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1
          Label operation: Swap 299888
          Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0x8001
          Label-switched-path ge-0/0/1.0:BypassLSP->192.168.0.1
          Label operation: Swap 299888, Push 299872(top)
          Label TTL action: prop-ttl, prop-ttl(top)
          State: < Active Int AckRequest>
          Age: 20          Metric: 1
          Validation State: unverified
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
          FECs bound to route: P2MP root-addr 192.168.0.5, lsp-id 99

```

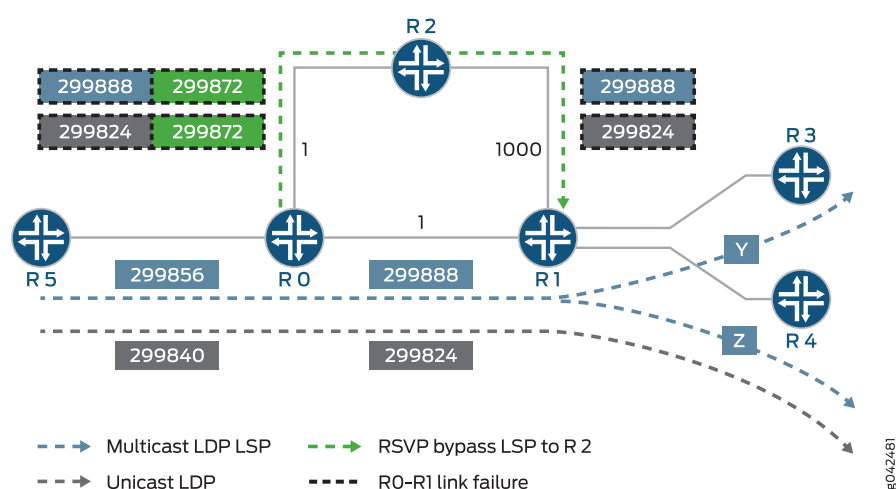
The main path for both unicast and multicast LDP LSP is through interface ge-0/0/1 (the link to Router R1), and the LFA path is through interface ge-0/0/2 (the link to Router R2). The LFA path is not used unless the ge-0/0/1 interface goes down.

Label 299840 is traffic arriving at Router R0 that corresponds to unicast LDP traffic to Router R4. Label 299856 is traffic arriving at Router 0 that corresponds to multicast LDP traffic from the root node R5 to the leaf egress nodes, R3 and R4.

As seen in the **show route detail** command output, the label operations for the protection path are the same for unicast LDP and multicast LDP traffic. The incoming LDP label at Router R0 is swapped to the LDP label advertised by Router R1 to Router R0. The RSVP label 299872 for the bypass tunnel is then pushed onto the packet. Penultimate hop-popping is used on the bypass tunnel, causing Router R2 to pop that label. Thus the packet arrives at Router R1 with the LDP label that it had originally advertised to Router R0.

Figure 82 on page 998 illustrates the label operation for unicast LDP and multicast LDP traffic protected by the RSVP bypass tunnel. The grey and blue boxes represent label values used under normal conditions for unicast and multicast LDP traffic, respectively. The dotted boxes represent label values used when the R0-R1 link fails.

Figure 82: LDP Link Protection Label Operation



Case C: Dynamic RSVP LSP Only

In this mode of operation, LFA is not used at all. A dynamic RSVP bypass LSP is automatically created in order to provide link protection. The output from the **show route detail** command and the label operations are similar to Case B, LFA and dynamic RSVP LSP mode.

Sample Multicast LDP Link Protection Configuration

To enable multicast LDP link protection, the following configuration is required on Router R0:

NOTE: In this sample, multicast LDP link protection is enabled on the ge-1/0/0 interface of Router R0 that connects to Router R1, although typically all the interfaces need to be configured for link protection.

Router R0


```

protocols {
  rsvp {
    interface all;
    interface ge-0/0/0.0 {
      disable;
    }
  }
  mpls {
    interface all;
    interface ge-0/0/0.0 {
      disable;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface lo0.0;
      interface ge-0/0/1.0 {
        link-protection;
      }
      interface ge-0/0/2.0;
      interface ge-0/0/3.0;
    }
  }
  ldp {
    make-before-break {
      timeout seconds;
      switchover-delay seconds;
    }
    interface ge-1/1/0.0 {
      link-protection {
        disable;
        dynamic-rsvp-lsp;
      }
    }
  }
}

```

The following configuration statements apply to the different modes of multicast LDP protection as follows:

- **link-protection** statement at [edit protocols ospf interface ge-0/0/1.0]

This configuration is applied only for Case A (LFA only) and Case B (LFA and dynamic RSVP LSP) modes of multicast LDP link protection. Configuring link protection under an IGP is not required for Case C (dynamic RSVP LSP only).

- **link-protection** statement at **[edit protocols ldp interface ge-0/0/1.0]**

This configuration is required for all modes of multicast LDP protection. However, if the only LDP traffic present is unicast, and dynamic RSVP bypasses are not required, then this configuration is not required, as the **link-protection** statement at the **[edit protocols ospf interface ge-0/0/1.0]** hierarchy level results in LFA action for the LDP unicast traffic.

- **dynamic-rsvp-lsp** statement at **[edit protocols ldp interface ge-0/0/1.0 link-protection]**

This configuration is applied only for Case B (LFA and dynamic RSVP LSP) and Case C (dynamic RSVP LSP only) modes of LDP link protection. Dynamic RSVP LSP configuration does not apply to Case A (LFA only).

Make-Before-Break

The make-before-break feature is enabled by default on Junos OS and provides some benefits for point-to-multipoint LSPs.

For a point-to-multipoint LSP, a label-switched router (LSR) selects the LSR that is its next hop to the root of the LSP as its upstream LSR. When the best path to reach the root changes, the LSR chooses a new upstream LSR. During this period, the LSP might be temporarily broken, resulting in packet loss until the LSP reconverges to a new upstream LSR. The goal of make-before-break in this case is to minimize the packet loss. In cases where the best path from the LSR to the root changes but the LSP continues to forward traffic to the previous next hop to the root, a new LSP should be established before the old LSP is withdrawn to minimize the duration of packet loss.

Taking for example, after a link failure, a downstream LSR (for instance, LSR-D) still receives and/or forwards packets to the other downstream LSRs, as it continues to receive packets from the one hop RSVP LSP. Once routing converges, LSR-D selects a new upstream LSR (LSR-U) for this point-to-multipoint LSP's FEC (FEC-A). The new LSR might already be forwarding packets for FEC-A to the downstream LSRs other than LSR-D. After LSR-U receives a label for FEC-A from LSR-D, it notifies LSR-D when it has learnt that LSP for FEC-A has been established from the root to itself. When LSR-D receives such a notification, it changes its next hop for the LSP root to LSR-U. This is a route delete and add operation on LSR-D. At this point, LSR-D does an LSP switchover, and traffic tunneled through RSVP LSP or LFA is dropped, and traffic from LSR-U is accepted. The new transit route for LSR-U is added. The RPF check is changed to accept traffic from LSR-U and to drop traffic from the old upstream LSR, or the old route is deleted and the new route is added.

The assumption is that LSR-U has received a make-before-break notification from its upstream router for the FEC-A point-to-multipoint LSP and has installed a forwarding state for the LSP. At that point it should signal LSR-D by means of make-before-break notification that it has become part of the tree identified by FEC-A and that LSR-D should initiate its switchover to the LSP. Otherwise, LSR-U should remember that it needs to send notification to LSR-D when it receives a make-before-break notification from the upstream

LSR for FEC-A and installs a forwarding state for this LSP. LSR-D continues to receive traffic from the old next hop to the root node using one hop RSVP LSP or LFA path until it switches over to the new point-to-multipoint LSP to LSR-U.

The make-before-break functionality with multicast LDP link protection includes the following features:

- **Make-before-break capability**

An LSR advertises that it is capable of handling make-before-break LSPs using the capability advertisement. If the peer is not make-before-break capable, the make-before-break parameters are not sent to this peer. If an LSR receives a make-before-break parameter from a downstream LSR (LSR-D) but the upstream LSR (LSR-U) is not make-before-break capable, the LSR immediately sends a make-before-break notification to LSR-D, and the make-before-break capable LSP is not established. Instead, the normal LSP is established.

- **Make-before-break status code**

The make-before-break status code includes:

- 1—make-before-break request
- 2—make-before-break acknowledgment

When a downstream LSR sends a label-mapping message for point-to-multipoint LSP, it includes the make-before-break status code as 1 (request). When the upstream LSR updates the forwarding state for the point-to-multipoint LSP, it informs the downstream LSR with a notification message containing the make-before-break status code as 2 (acknowledgment). At that point, the downstream LSR does an LSP switchover.

Caveats and Limitations

The Junos OS implementation of the LDP link protection feature has the following caveats and limitations:

- Make-before-break is not supported for the following point-to-multipoint LSPs on an egress LSR:
 - Next-generation multicast virtual private network (MVPN) with virtual routing and forwarding (VRF) label
 - Static LSP
- The following features are not supported:
 - Nonstop active routing for point-to-multipoint LSP in Junos OS Releases 12.3, 13.1 and 13.2
 - Graceful restart switchover point-to-multipoint LSP
 - Link protection for routing instance

SEE ALSO

Example: Configuring LDP Link Protection

IN THIS SECTION

- [Requirements](#) | 1002
- [Overview](#) | 1002
- [Configuration](#) | 1004
- [Verification](#) | 1011

This example shows how to configure Label Distribution Protocol (LDP) link protection for both unicast and multicast LDP label-switched paths (LSPs).

Requirements

This example uses the following hardware and software components:

- Six routers that can be a combination of M Series, MX Series, or T Series routers with one root node and two leaf nodes running a point-to-multipoint LDP LSP.
- Junos OS Release 12.3 or later running on all the routers.

Before you begin:

1. Configure the device interfaces.
2. Configure the following protocols:
 - a. RSVP
 - b. MPLS
 - c. OSPF or any other IGP
 - d. LDP

Overview

LDP link protection enables fast reroute of traffic carried over LDP LSPs in case of a link failure. LDP point-to-multipoint LSPs can be used to send traffic from a single root or ingress node to a number of leaf nodes or egress nodes traversing one or more transit nodes. When one of the links of the point-to-multipoint tree fails, the subtrees can get detached until the IGP reconverges and multicast LDP initiates label mapping using the best path from the downstream router to the new upstream router. To protect the traffic in the event of a link failure, you can configure an explicit tunnel so that traffic can be rerouted using the tunnel. Junos OS supports make-before-break capabilities to ensure minimum packet loss when attempting to

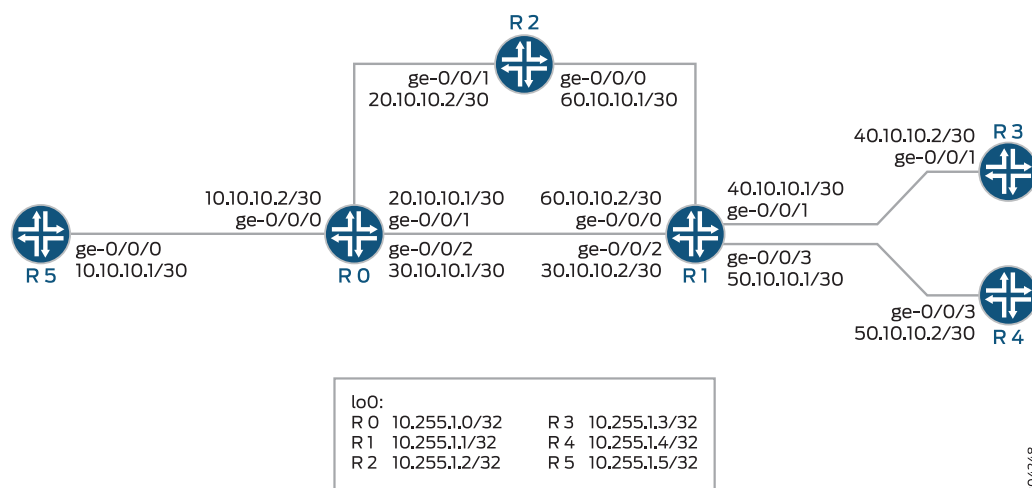
signal a new LSP path before tearing down the old LSP path. This feature also adds targeted LDP support for multicast LDP link protection.

When configuring LDP link protection, be aware of the following considerations:

- Configure traffic engineering under IGP (if it is not supported by default), and include the interfaces configured for MPLS and RSVP so that constrained-based link protected dynamic RSVP LSP is signaled by RSVP using Constrained Shortest Path First (CSPF). When this condition is not satisfied, RSVP LSP might not come up and LDP cannot use it as a protected next hop.
- Configure a path between two label-switched routers (LSRs) to provide IP connectivity between the routers when there is a link failure. This enables CSPF to calculate an alternate path for link protection. When the connectivity between the routers is lost, the LDP targeted adjacency does not come up and dynamic RSVP LSP cannot be signaled, resulting in no protection for the LDP forwarding equivalence class (FEC) for which the peer is the downstream LSR.
- If link protection is active only on one LSR, then the other LSR should not be configured with the **strict-targeted-hellos** statement. This enables the LSR without link protection to allow asymmetric remote neighbor discovery and send periodic targeted hellos to the LSR that initiated the remote neighbor. When this condition is not satisfied, LDP targeted adjacency is not formed.
- LDP must be enabled on the loopback interface of the LSR to create remote neighbors based on LDP tunneling, LDP-based virtual private LAN service (VPLS), Layer 2 circuits, or LDP session protection. When this condition is not satisfied, LDP targeted adjacency is not formed.
- For unicast LDP LSP, loop-free alternate (LFA) should be configured in IGP.
- The ingress route to merge point should have at least one next hop avoiding the primary link between the merge point and the point of local repair for unicast LDP LSP.
- Point of local repair should have a unicast LDP label for the backup next hop to reach the merge point.

Topology

Figure 83: LDP Link Protection



In this example, Router R5 is the root connecting to two leaf nodes, Routers R3 and R4. Router R0 is the point of local repair.

Configuration

CLI Quick Configuration

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

R5

```
set interfaces ge-0/0/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.5/32
set routing-options router-id 10.255.1.5
set routing-options autonomous-system 100
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all metric 1
```



```

set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all link-protection dynamic-rsvp-lsp
set protocols ldp interface fxp0.0 disable
set protocols ldp p2mp

```

R0

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.10.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 20.10.10.1/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 30.10.10.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.0/32
set routing-options router-id 10.255.1.0
set routing-options autonomous-system 100
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all metric 1
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all link-protection dynamic-rsvp-lsp
set protocols ldp interface fxp0.0 disable
set protocols ldp p2mp

```

R1

```

set interfaces ge-0/0/0 unit 0 family inet address 60.10.10.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 40.10.10.1/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 30.10.10.2/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 50.10.10.1/30

```



```

set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.1/32
set routing-options router-id 10.255.1.1
set routing-options autonomous-system 100
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all metric 1
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all link-protection dynamic-rsvp-lsp
set protocols ldp interface fxp0.0 disable
set protocols ldp p2mp

```

R2

```

set interfaces ge-0/0/0 unit 0 family inet address 60.10.10.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 20.10.10.2/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.2/32
set routing-options router-id 10.255.1.2
set routing-options autonomous-system 100
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all link-protection dynamic-rsvp-lsp
set protocols ldp interface fxp0.0 disable
set protocols ldp p2mp

```

R3


```
set interfaces ge-0/0/1 unit 0 family inet address 40.10.10.2/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.3/32
set routing-options router-id 10.255.1.3
set routing-options autonomous-system 100
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all metric 1
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all link-protection dynamic-rsvp-lsp
set protocols ldp interface fxp0.0 disable
set protocols ldp p2mp root-address 10.255.1.5 lsp-id 1
```

R4

```
set interfaces ge-0/0/3 unit 0 family inet address 50.10.10.2/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.1.4/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all metric 1
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all link-protection dynamic-rsvp-lsp
set protocols ldp interface fxp0.0 disable
set protocols ldp p2mp root-address 10.255.1.5 lsp-id 1
```

Step-by-Step Procedure

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

To configure Router R0:

1. Configure the Router R0 interfaces.

```
[edit interfaces]
user@R0# set ge-0/0/0 unit 0 family inet address 10.10.10.2/30
user@R0# set ge-0/0/0 unit 0 family mpls
user@R0# set ge-0/0/1 unit 0 family inet address 20.10.10.1/30
user@R0# set ge-0/0/1 unit 0 family mpls
user@R0# set ge-0/0/2 unit 0 family inet address 30.10.10.1/30
user@R0# set ge-0/0/2 unit 0 family mpls
user@R0# set lo0 unit 0 family inet address 10.255.1.0/32
```

2. Configure the router ID and autonomous system of Router R0.

```
[edit routing-options]
user@R0# set router-id 10.255.1.0
user@R0# set autonomous-system 100
```

3. Enable RSVP on all the interfaces of Router R0 (excluding the management interface).

```
[edit protocols]
user@R0# set rsvp interface all
user@R0# set rsvp interface fxp0.0 disable
```

4. Enable MPLS on all the interfaces of Router R0 (excluding the management interface) along with traffic engineering capabilities.

```
[edit protocols]
user@R0# set mpls traffic-engineering
user@R0# set mpls interface all
user@R0# set mpls interface fxp0.0 disable
```

5. Enable OSPF on all the interfaces of Router R0 (excluding the management interface), assign equal cost metric for the links, and enable traffic engineering capabilities.

```
[edit protocols]
user@R0# set ospf traffic-engineering
```



```
user@R0# set ospf area 0.0.0.0 interface all metric 1
user@R0# set ospf area 0.0.0.0 interface fxp0.0 disable
```

NOTE: For multicast LDP link protection with loop-free alternative (LFA), enable the following configuration under the **[edit protocols]** hierarchy level:

```
set ospf area 0 interface all link-protection
```

6. Enable LDP on all the interfaces of Router R0 (excluding the management interface) and configure link protection with dynamic RSVP bypass LSP.

```
[edit protocols]
user@R0# set ldp interface all link-protection dynamic-rsvp-lsp
user@R0# set ldp interface fxp0.0 disable
user@R0# set ldp p2mp
```

Results

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

```
user@R0# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.10.2/30;
    }
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 20.10.10.1/30;
    }
    family mpls;
  }
}
```



```

ge-0/0/2 {
  unit 0 {
    family inet {
      address 30.10.10.1/30;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.1.0/32;
    }
  }
}

```

```

user@R0# show routing-options
router-id 10.255.1.0;
autonomous-system 100;

```

```

user@R0# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  traffic-engineering;
  interface all;
  interface fxp0.0 {
    disable;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all {
      metric 1;
    }
    interface fxp0.0 {
      disable;
    }
  }
}

```



```

    }
  }
  ldp {
    interface all {
      link-protection {
        dynamic-rsvp-lsp;
      }
    }
    interface fxp0.0 {
      disable;
    }
    p2mp;
  }
}

```

Verification

IN THIS SECTION

- [Verifying the Bypass RSVP LSP Path | 1011](#)
- [Verifying Label Operation | 1012](#)

Verify that the configuration is working properly.

Verifying the Bypass RSVP LSP Path

Purpose

Verify that the bypass RSVP LSP path has been created on the point of local repair (PLR).

Action

From operational mode, run the **show route tale mpls.0** command.

user@R0> **show route tale mpls.0**

```

mpls.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 05:28:13, metric 1
            Receive
1          *[MPLS/0] 05:28:13, metric 1
            Receive

```



```

2          *[MPLS/0] 05:28:13, metric 1
          Receive
13         *[MPLS/0] 05:28:13, metric 1
          Receive
299792     *[LDP/9] 00:41:41, metric 1
          > to 30.10.10.2 via ge-0/0/2.0, Pop
299792(S=0) *[LDP/9] 00:41:41, metric 1
          > to 30.10.10.2 via ge-0/0/2.0, Pop
299808     *[LDP/9] 00:41:41, metric 1
          > to 20.10.10.2 via ge-0/0/1.0, Pop
299808(S=0) *[LDP/9] 00:41:41, metric 1
          > to 20.10.10.2 via ge-0/0/1.0, Pop
299920     *[RSVP/7/1] 01:51:43, metric 1
          > to 30.10.10.2 via ge-0/0/2.0, label-switched-path
ge-0/0/0.0:BypassLSP->10.255.1.1
299920(S=0) *[RSVP/7/1] 01:51:43, metric 1
          > to 30.10.10.2 via ge-0/0/2.0, label-switched-path
ge-0/0/0.0:BypassLSP->10.255.1.1
299936     *[RSVP/7/1] 01:51:25, metric 1
          > to 20.10.10.2 via ge-0/0/1.0, label-switched-path
ge-0/0/0.0:BypassLSP->10.255.1.2
299936(S=0) *[RSVP/7/1] 01:51:25, metric 1
          > to 20.10.10.2 via ge-0/0/1.0, label-switched-path
ge-0/0/0.0:BypassLSP->10.255.1.2
299952     *[LDP/9] 00:06:11, metric 1
          > to 10.10.10.1 via ge-0/0/0.0, Pop
299952(S=0) *[LDP/9] 00:06:11, metric 1
          > to 10.10.10.1 via ge-0/0/0.0, Pop
299968     *[LDP/9] 00:05:39, metric 1
          > to 30.10.10.2 via ge-0/0/2.0, Swap 299984
299984     *[LDP/9] 00:05:38, metric 1
          > to 30.10.10.2 via ge-0/0/2.0, Swap 300000
300000     *[LDP/9] 00:05:15, metric 1
          > to 30.10.10.2 via ge-0/0/2.0, Swap 300016

```

Meaning

When the R0-R1 link goes down, the RSVP bypass LSP is used to route traffic.

Verifying Label Operation

Purpose

Verify the label swapping at the PLR.

Action

From operational mode, run the **show route table mpls.0 label label extensive** command.

user@R0> **show route table mpls.0 label 300000 extensive**

```
mpls.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
300000 (1 entry, 1 announced)
TSI:
KRT in-kernel 300000 /52 -> {Swap 300016}
    *LDP      Preference: 9
              Next hop type: Router, Next hop index: 589
              Address: 0x9981610
              Next-hop reference count: 2
              Next hop: 30.10.10.2 via ge-0/0/2.0, selected
              Label operation: Swap 300016
              Load balance label: Label 300016: None;
              Session Id: 0x2
              State: <Active Int>
              Local AS: 100
              Age: 12:50      Metric: 1
              Validation State: unverified
              Task: LDP
              Announcement bits (1): 1-KRT
              AS path: I
              Prefixes bound to route: 10.255.1.4/32
```

Meaning

The label is bound to reach Router R4, which is a leaf node.

SEE ALSO

| [LDP Link Protection Overview](#) | 983

Understanding Multicast-Only Fast Reroute

Starting in Junos OS Release 14.1, Multicast-only fast reroute (MoFRR) functionality is available, in which packet loss is minimized in PIM and multipoint LDP domains. MoFRR minimizes packet loss in a network when there is a link failure. It works by enhancing multicast routing protocols like Protocol Independent Multicast (PIM) and multipoint Label Distribution Protocol (multipoint LDP). MoFRR is supported on MX Series routers with MPC line cards. As a prerequisite, the router must be set to **network-services enhanced-ip** mode, and all the line-cards in the router must be MPCs.

With MoFRR enabled, join messages are sent on primary and backup upstream paths. Data packets are received from both the primary path and the backup paths. The redundant packets are discarded based on priority (weights that are assigned to the primary and backup paths). When a failure is detected on the primary path, the repair is made by changing the interface on which packets are accepted to the secondary interface. Because the repair is local, it is fast—greatly improving convergence times in the event of a link failure on the primary path.

Currently, the most likely real-world application for MoFRR is streaming IPTV. IPTV streams are multicast as UDP streams. Therefore, any lost packets are not retransmitted, and this can result in a less-than-satisfactory user experience. MoFRR can be used to improve this situation.

When fast reroute is applied to unicast streams, an upstream router preestablishes MPLS label-switched paths (LSPs) or precomputes an IP loop-free alternate (LFA) fast reroute backup path to handle failure of a segment in the downstream path.

In multicast routing, the traffic distribution graphs are usually originated by the receiver. This is unlike unicast routing, which usually establishes the path from the source to the receiver. Protocols that are capable of establishing multicast distribution graphs are PIM (for IP), multipoint LDP (for MPLS), and RSVP-TE (for MPLS). Of these, PIM and multipoint LDP receivers initiate the distribution graph setup, and therefore these are the two multicast protocols for which MoFRR is supported.

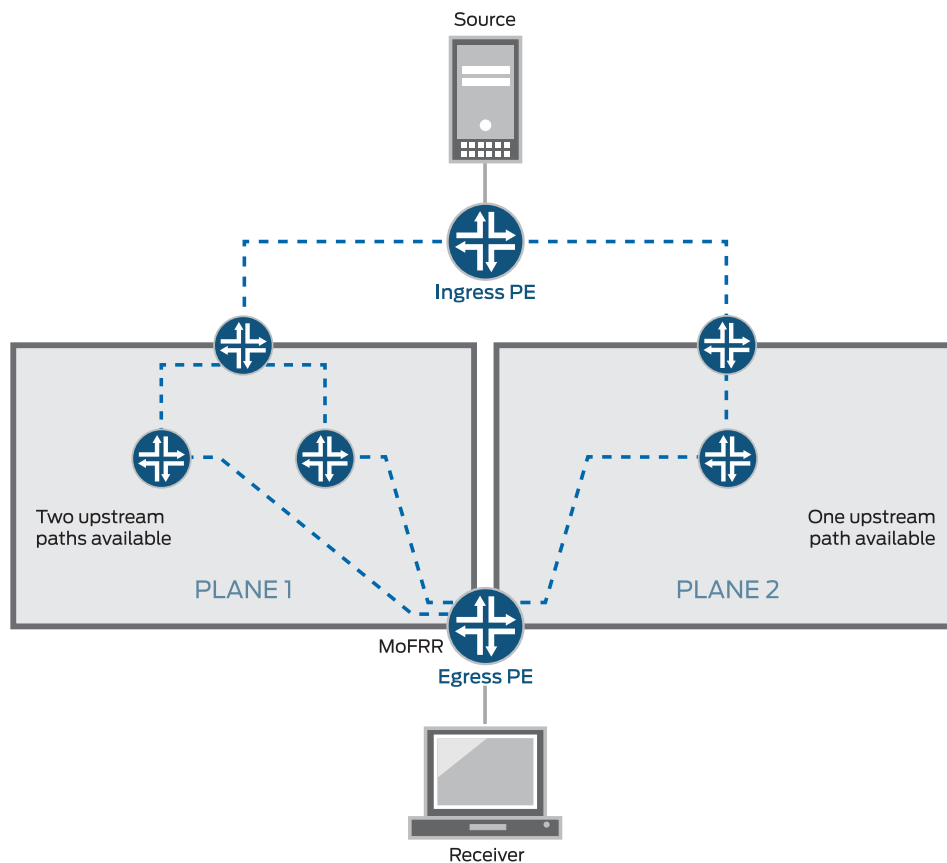
In a multicast tree, performing a reactive repair upon detection of a network-component failure can lead to significant traffic loss due to delay in setting up the alternative path. MoFRR reduces traffic loss in a multicast distribution tree when a network component fails. With MoFRR, one of the downstream routers that supports this feature sets up an alternative path toward the source to receive a backup live stream of the same multicast traffic. When a failure is detected on the primary stream, the MoFRR router switches to the backup stream.

With MoFRR enabled, for each (S,G) entry, two of the available upstream interfaces are used to send a join message and to receive multicast traffic. The protocol attempts to select two disjoint paths if two such paths are available. If disjoint paths are not available, the protocol selects two non-disjoint paths. If two non-disjoint paths are not available, only a primary path is selected with no backup. MoFRR is supported for both IPv4 and IPv6 protocol families.

In the context of load balancing, MoFRR prioritizes the disjoint backup in favor of load balancing the available paths.

Figure 84 on page 1015 shows two paths from the egress provider edge (PE) router to the ingress PE router.

Figure 84: MoFRR Sample Topology



When enabled with MoFRR functionality, the egress router sets up two multicast trees, a primary path and a backup path, toward the multicast source for each (S,G). In other words, the egress router propagates the same (S,G) join messages toward two different upstream neighbors, thus creating two multicast trees.

One of the multicast trees goes through plane 1 and the other through plane 2, as shown in Figure 84 on page 1015. For each (S,G), the egress PE router forwards traffic received on the primary path and drops traffic received on the backup path.

MoFRR is supported on both equal-cost multipath (ECMP) paths and non-ECMP paths. Unicast loop-free alternate (LFA) routes need to be enabled to support MoFRR on non-ECMP paths. LFA routes are enabled with the **link-protection** statement in the interior gateway protocol (IGP) configuration. When you enable link protection on an OSPF or IS-IS interface, Junos OS creates a backup LFA path to the primary next hop for all destination routes that traverse the protected interface.

Junos OS implements MoFRR in the IP network for IP MoFRR and at the MPLS label-edge router (LER) for multipoint LDP MoFRR.

Multipoint LDP MoFRR is used at the egress node of an MPLS network, where the packets are forwarded to an IP network. In the case of multipoint LDP MoFRR, the two paths toward the upstream PE router are established for receiving two streams of MPLS packets at the LER. One of the streams (the primary) is accepted, and the other one (the backup) is dropped at the LER. The backup stream is accepted if the primary path fails. A prerequisite for this feature is inband signaling support, as described in [“Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs” on page 1072](#).

PIM Functionality

Junos OS supports MoFRR for shortest-path tree (SPT) joins in PIM source-specific multicast (SSM) and any-source multicast (ASM). MoFRR is supported for both SSM and ASM ranges. To enable MoFRR for (*,G) joins, the `mofrr-asm-starg` configuration statement needs to be included. For each group G, either (S,G) or (*,G) (not both) will undergo MoFRR. (S,G) always takes precedence over (*,G).

With MoFRR enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream RPF next hops with two (primary and backup) interfaces.

When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.

MoFRR can be enabled along with PIM join load balancing (with the `join-load-balance automatic` statement). However, in such cases the distribution of join messages among the links might not be even. When a new ECMP link is added, join messages on the primary path are redistributed and load-balanced. The join messages on the backup path might still follow the same path and might not be evenly redistributed.

MoFRR is enabled with a `[edit routing-options multicast stream-protection]` configuration and is managed by a set of filter policies. When an egress PIM router receives a join message or an IGMP report, the router checks for the MoFRR configuration.

If the MoFRR configuration is not present, PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 84 on page 1015](#)).

If the MoFRR configuration is present, Junos OS checks for a policy configuration.

If a policy is not present, Junos OS checks for primary and backup paths (upstream interfaces), and takes the following actions:

- If primary and backup paths are not available—PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 84 on page 1015](#)).
- If primary and backup paths are available—PIM sends the join message upstream toward two of the available upstream neighbors. Junos OS sets up primary and secondary multicast paths to receive multicast traffic (for example, plane 1 in [Figure 84 on page 1015](#)).

If a policy is present, Junos OS checks whether the policy allows MoFRR for this (S,G), and takes the following actions:

- If the policy check fails—PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 84 on page 1015](#)).
- If the policy check passes—Junos OS checks for primary and backup paths (upstream interfaces).
 - If the primary and backup paths are not available, PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 84 on page 1015](#)).
 - If the primary and backup paths are available, PIM sends the join message upstream toward two of the available upstream neighbors. Junos OS sets up primary and secondary multicast paths to receive multicast traffic (for example, plane 1 in [Figure 84 on page 1015](#)).

Multipoint LDP Functionality

To avoid MPLS traffic duplication, the usual implementation of multipoint LDP selects only one upstream path. (See section 2.4.1.1. Determining One's 'upstream LSR' in RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*.)

For multipoint LDP MoFRR, the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other path becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop.

A forwarding equivalency class (FEC) is a group of IP packets that are forwarded in the same manner, over the same path, and with the same forwarding treatment. Normally, the label that is put on a particular packet represents the FEC to which that packet is assigned. In MoFRR, two routes are placed into the mpls.0 table for each FEC—one route for the primary label and the other route for the backup label.

If there are parallel links toward the same immediate upstream node, both parallel links are considered to be the primary. At any point in time, the upstream node sends traffic on only one of the multiple parallel links.

A bud node is an LSR that is an egress LSR, but also has one or more directly connected downstream LSRs. In the case of a bud node, the traffic from the primary upstream path is forwarded to a downstream LSR. If the primary upstream path fails, the MPLS traffic from the backup upstream path is forwarded to the downstream LSR. This means that the downstream LSR next hop is added to both MPLS routes along with the egress next hop.

MoFRR for multipoint LDP is enabled with a **[edit routing-options multicast stream-protection]** configuration and is managed by a set of filter policies.

If the multipoint LDP point-to-multipoint FEC is enabled for MoFRR, the following additional considerations are factored into upstream path selection:

- The targeted LDP sessions are skipped if there is a nontargeted LDP session. If there is a single targeted LDP session, the targeted LDP session is selected, but the corresponding point-to-multipoint FEC loses the MoFRR capability because there is no interface associated with the targeted LDP session.
- All interfaces that belong to the same upstream LSR are considered to be the primary path.
- For any root-node route updates, the upstream path is changed based on the latest next hops from the IGP. If a better path is available, multipoint LDP attempts to switch to the better path.

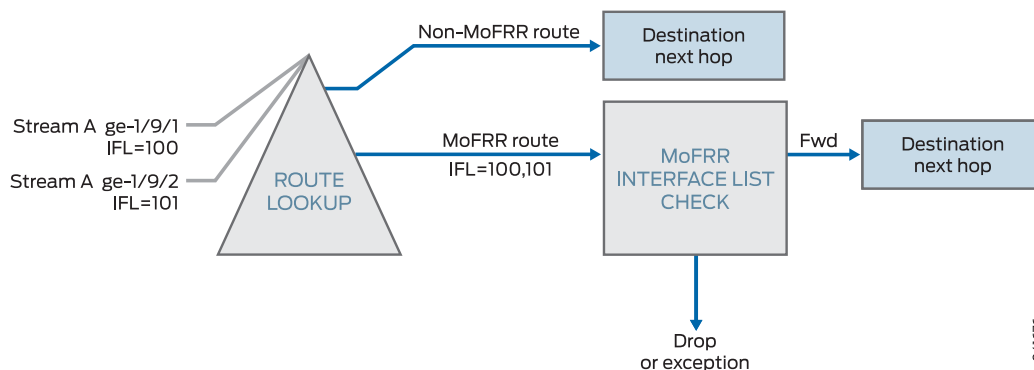
Packet Forwarding

For both PIM and multipoint LDP, multicast source stream selection is performed at the incoming interface. This prevents duplicate streams from being sent across the fabric and prevents multiple route lookups that result in drops, thus preserving fabric bandwidth and maximizing forwarding performance.

For PIM, each IP multicast stream contains the same destination address. Regardless of the interface on which the packets arrive, the packets have the same route. An interface list is attached to the route. Junos OS checks the interface upon which each packet arrives and forwards only those that are from the primary interface. If the interface matches a secondary interface, the packets are dropped. If no match is found, the packets are handled as exceptions in the control plane.

This process is shown in [Figure 85 on page 1018](#).

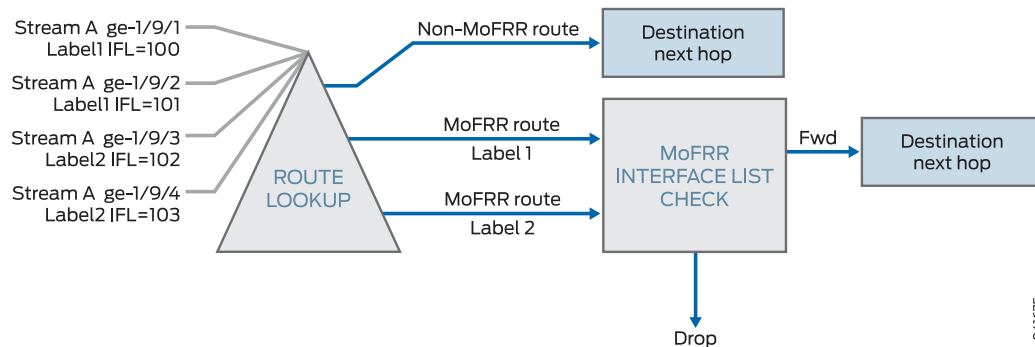
Figure 85: MoFRR IP Route Lookup in the Packet Forwarding Engine



For multipoint LDP, multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.

This process is shown in [Figure 86 on page 1019](#).

Figure 86: MoFRR MPLS Route Lookup in the Packet Forwarding Engine



Limitations and Caveats

MoFRR has the following limitations and caveats:

- MoFRR failure detection is supported for immediate link protection of the router on which MoFRR is enabled and not on all the links (end-to-end) in the multicast traffic path.
- MoFRR supports FRR on two selected disjoint paths toward the source. Two of the selected upstream neighbors cannot be on the same interface—in other words, two upstream neighbors on a LAN segment. The same is true if the upstream interface happens to be a multicast tunnel interface.
- Detection of the maximum end-to-end disjoint upstream paths is not supported. The egress router only makes sure that there is a disjoint upstream node (the immediate previous hop). PIM and multipoint LDP do not support the equivalent of explicit route objects (EROs). Hence, disjoint upstream path detection is limited to control over the immediately previous hop node. Because of this limitation, the path to the upstream node of the previous hop selected as primary and backup might be shared.
- MoFRR does not apply to multipoint LDP traffic received on an RSVP tunnel because the RSVP tunnel is not associated with any interface.
- Some traffic loss is seen in the following scenarios:
 - A better upstream path becomes available on an egress node.
 - MoFRR is enabled or disabled on the egress node while there is an active traffic stream flowing.
- PIM join load balancing for join messages for backup paths are not supported.
- For a multicast group G, MoFRR is not allowed for both (S,G) and (*,G) join messages. (S,G) join messages have precedence over (*,G).
- MoFRR is not supported for multicast traffic streams that use two different multicast groups. Each (S,G) combination is treated as a unique multicast traffic stream.
- The bidirectional PIM range is not supported for MoFRR.

- PIM dense-mode is not supported for MoFRR
- Mixed upstream MoFRR is not supported. This refers to PIM multipoint LDP in-band signaling, wherein one upstream path is through multipoint LDP and the second upstream path is through PIM.
- Multicast statistics for the backup traffic stream are not maintained by PIM and therefore are not available in the operational output of **show** commands.
- Multipoint LDP labels as inner labels are not supported.
- If the source is reachable through multiple ingress provider edge (PE) routers, multipoint LDP MoFRR is not supported.
- Targeted upstream sessions are not selected as the upstream node for MoFRR.
- Rate monitoring is not supported.
- Multipoint LDP link protection on the backup path is not supported because there is no support for MoFRR inner labels.

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, Multicast-only fast reroute (MoFRR) functionality is available, in which packet loss is minimized in PIM and multipoint LDP domains.

RELATED DOCUMENTATION

Configuring Multicast-Only Fast Reroute 1021
<i>Example: Configuring Multicast-Only Fast Reroute in a PIM Domain</i>
Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain 1024

Configuring Multicast-Only Fast Reroute

You can configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure.

When fast reroute is applied to unicast streams, an upstream router preestablishes MPLS label-switched paths (LSPs) or precomputes an IP loop-free alternate (LFA) fast reroute backup path to handle failure of a segment in the downstream path.

In multicast routing, the traffic distribution graphs are usually originated by the receiver. This is unlike unicast routing, which usually establishes the path from the source to the receiver. Protocols that are capable of establishing multicast distribution graphs are PIM (for IP), multipoint LDP (for MPLS) and RSVP-TE (for MPLS). Of these, PIM and multipoint LDP receivers initiate the distribution graph setup, and therefore:

- On the QFX series, MoFRR is supported in PIM domains.
- On the MX Series and SRX Series, MoFRR is supported in PIM and multipoint LDP domains.

The configuration steps are the same for enabling MoFRR for PIM on all devices that support this feature, unless otherwise indicated. Configuration steps that are not applicable to multipoint LDP MoFRR are also indicated.

(For MX Series routers only) MoFRR is supported on MX Series routers with MPC line cards. As a prerequisite, all the line cards in the router must be MPCs.

To configure MoFRR on routers or switches:

1. (For MX Series and SRX Series routers only) Set the router to enhanced IP mode.

```
[edit chassis]
user@host# set network-services enhanced-ip
```

2. Enable MoFRR.

```
[edit routing-options multicast]
user@host# set stream-protection
```

3. (Optional) Configure a routing policy that filters for a restricted set of multicast streams to be affected by your MoFRR configuration.

You can apply filters that are based on source or group addresses.

For example:

```
[edit policy-options]
policy-statement mofrr-select {
  term A {
    from {
      source-address-filter 225.1.1.1/32 exact;
    }
    then {
      accept;
    }
  }
  term B {
    from {
      source-address-filter 226.0.0.0/8 orlonger;
    }
    then {
      accept;
    }
  }
  term C {
    from {
      source-address-filter 227.1.1.0/24 orlonger;
      source-address-filter 227.4.1.0/24 orlonger;
      source-address-filter 227.16.1.0/24 orlonger;
    }
    then {
      accept;
    }
  }
  term D {
    from {
      source-address-filter 227.1.1.1/32 exact
    }
    then {
      reject; #MoFRR disabled
    }
  }
  ...
}
```

4. (Optional) If you configured a routing policy to filter the set of multicast groups to be affected by your MoFRR configuration, apply the policy for MoFRR stream protection.


```
[edit routing-options multicast stream-protection]
user@host# set policy policy-name
```

For example:

```
routing-options {
  multicast {
    stream-protection {
      policy mofrr-select
    }
  }
}
```

5. (Optional) In a PIM domain with MoFRR, allow MoFRR to be applied to any-source multicast (ASM) (*,G) joins.

This is not supported for multipoint LDP MoFRR.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-asm-starg
```

6. (Optional) In a PIM domain with MoFRR, allow only a disjoint RPF (an RPF on a separate plane) to be selected as the backup RPF path.

This is not supported for multipoint LDP MoFRR. In a multipoint LDP MoFRR domain, the same label is shared between parallel links to the same upstream neighbor. This is not the case in a PIM domain, where each link forms a neighbor. The **mofrr-disjoint-upstream-only** statement does not allow a backup RPF path to be selected if the path goes to the same upstream neighbor as that of the primary RPF path. This ensures that MoFRR is triggered only on a topology that has multiple RPF upstream neighbors.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-disjoint-upstream-only
```

7. (Optional) In a PIM domain with MoFRR, prevent sending join messages on the backup path, but retain all other MoFRR functionality.

This is not supported for multipoint LDP MoFRR.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-no-backup-join
```


8. (Optional) In a PIM domain with MoFRR, allow new primary path selection to be based on the unicast gateway selection for the unicast route to the source and to change when there is a change in the unicast selection, rather than having the backup path be promoted as primary. This ensures that the primary RPF hop is always on the best path.

When you include the **mofrr-primary-selection-by-routing** statement, the backup path is not guaranteed to get promoted to be the new primary path when the primary path goes down.

This is not supported for multipoint LDP MoFRR.

```
[edit routing-options multicast stream-protection]
user@host# set mofrr-primary-path-selection-by-routing
```

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

Understanding Multicast-Only Fast Reroute on Switches

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

IN THIS SECTION

- [Requirements | 1025](#)
- [Overview | 1025](#)
- [CLI Quick Configuration | 1026](#)
- [Configuration | 1034](#)
- [Verification | 1041](#)

This example shows how to configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure.

Multipoint LDP MoFRR is used at the egress node of an MPLS network, where the packets are forwarded to an IP network. In the case of multipoint LDP MoFRR, the two paths toward the upstream provider edge (PE) router are established for receiving two streams of MPLS packets at the label-edge router (LER). One of the streams (the primary) is accepted, and the other one (the backup) is dropped at the LER. The backup stream is accepted if the primary path fails.

Requirements

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

In a multipoint LDP domain, for MoFRR to work, only the egress PE router needs to have MoFRR enabled. The other routers do not need to support MoFRR.

MoFRR is supported on MX Series platforms with MPC line cards. As a prerequisite, the router must be set to **network-services enhanced-ip** mode, and all the line-cards in the platform must be MPCs.

This example requires Junos OS Release 14.1 or later on the egress PE router.

Overview

In this example, Device R3 is the egress edge router. MoFRR is enabled on this device only.

OSPF is used for connectivity, though any interior gateway protocol (IGP) or static routes can be used.

For testing purposes, routers are used to simulate the source and the receiver. Device R4 and Device R8 are configured to statically join the desired group by using the **set protocols igmp interface interface-name static group group** command. In the case when a real multicast receiver host is not available, as in this example, this static IGMP configuration is useful. On the receivers, to make them listen to the multicast group address, this example uses **set protocols sap listen group**.

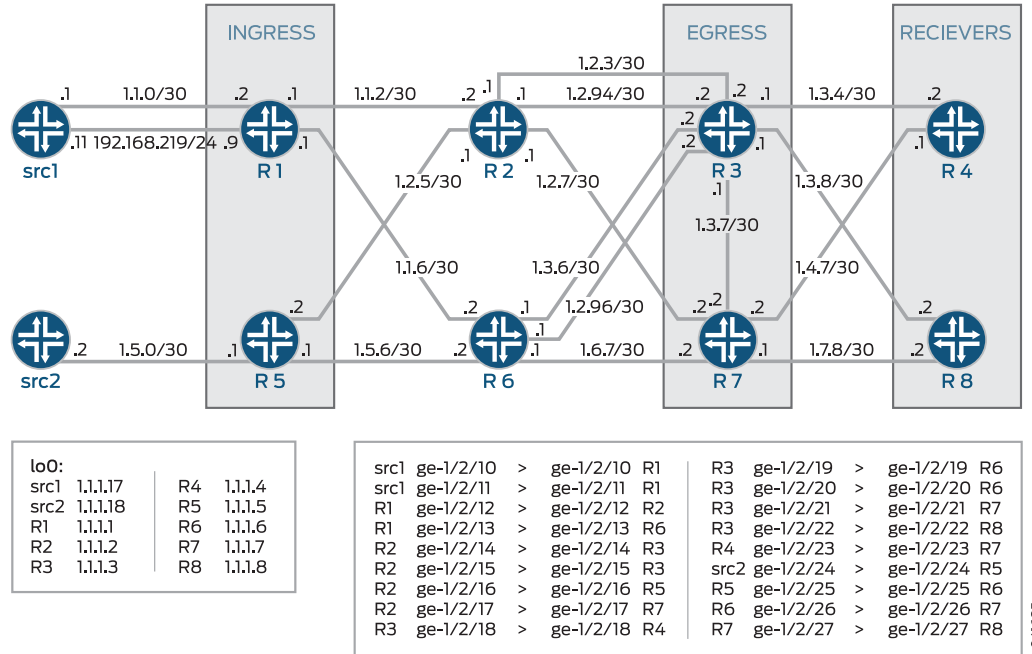
MoFRR configuration includes a policy option that is not shown in this example, but is explained separately. The option is configured as follows:

```
stream-protection {
  policy policy-name;
}
```

Topology

Figure 87 on page 1026 shows the sample network.

Figure 87: MoFRR in a Multipoint LDP Domain



"CLI Quick Configuration" on page 1026 shows the configuration for all of the devices in Figure 87 on page 1026.

The section "Configuration" on page 1034 describes the steps on Device R3.

CLI Quick Configuration

CLI Quick Configuration

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

Device src1

```
set interfaces ge-1/2/10 unit 0 description src1-to-R1
set interfaces ge-1/2/10 unit 0 family inet address 1.1.0.1/30
set interfaces ge-1/2/11 unit 0 description src1-to-R1
set interfaces ge-1/2/11 unit 0 family inet address 192.168.219.11/24
set interfaces lo0 unit 0 family inet address 1.1.1.17/32
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Device src2


```

set interfaces ge-1/2/24 unit 0 description src2-to-R5
set interfaces ge-1/2/24 unit 0 family inet address 1.5.0.2/30
set interfaces lo0 unit 0 family inet address 1.1.1.18/32
set protocols rsvp interface all
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Device R1

```

set interfaces ge-1/2/12 unit 0 description R1-to-R2
set interfaces ge-1/2/12 unit 0 family inet address 1.1.2.1/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/13 unit 0 description R1-to-R6
set interfaces ge-1/2/13 unit 0 family inet address 1.1.6.1/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/10 unit 0 description R1-to-src1
set interfaces ge-1/2/10 unit 0 family inet address 1.1.0.2/30
set interfaces ge-1/2/11 unit 0 description R1-to-src1
set interfaces ge-1/2/11 unit 0 family inet address 192.168.219.9/30
set interfaces lo0 unit 0 family inet address 1.1.1.1/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.1
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.3
set protocols bgp group ibgp neighbor 1.1.1.7
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/12.0
set protocols ldp interface ge-1/2/13.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim rp static address 1.1.1.5
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/10.0
set protocols pim interface ge-1/2/11.0

```



```

set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
  orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 1.1.1.7/32 orlonger
set policy-options policy-statement mldppim-ex term A from source-address-filter 1.1.0.0/30 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 10

```

Device R2

```

set interfaces ge-1/2/12 unit 0 description R2-to-R1
set interfaces ge-1/2/12 unit 0 family inet address 1.1.2.2/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/14 unit 0 description R2-to-R3
set interfaces ge-1/2/14 unit 0 family inet address 1.2.3.1/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces ge-1/2/16 unit 0 description R2-to-R5
set interfaces ge-1/2/16 unit 0 family inet address 1.2.5.1/30
set interfaces ge-1/2/16 unit 0 family mpls
set interfaces ge-1/2/17 unit 0 description R2-to-R7
set interfaces ge-1/2/17 unit 0 family inet address 1.2.7.1/30
set interfaces ge-1/2/17 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 description R2-to-R3
set interfaces ge-1/2/15 unit 0 family inet address 1.2.94.1/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.2/32
set interfaces lo0 unit 0 family mpls
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp

```



```

set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
  orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 10

```

Device R3

```

set chassis network-services enhanced-ip
set interfaces ge-1/2/14 unit 0 description R3-to-R2
set interfaces ge-1/2/14 unit 0 family inet address 1.2.3.2/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces ge-1/2/18 unit 0 description R3-to-R4
set interfaces ge-1/2/18 unit 0 family inet address 1.3.4.1/30
set interfaces ge-1/2/18 unit 0 family mpls
set interfaces ge-1/2/19 unit 0 description R3-to-R6
set interfaces ge-1/2/19 unit 0 family inet address 1.3.6.2/30
set interfaces ge-1/2/19 unit 0 family mpls
set interfaces ge-1/2/21 unit 0 description R3-to-R7
set interfaces ge-1/2/21 unit 0 family inet address 1.3.7.1/30
set interfaces ge-1/2/21 unit 0 family mpls
set interfaces ge-1/2/22 unit 0 description R3-to-R8
set interfaces ge-1/2/22 unit 0 family inet address 1.3.8.1/30
set interfaces ge-1/2/22 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 description R3-to-R2
set interfaces ge-1/2/15 unit 0 family inet address 1.2.94.2/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces ge-1/2/20 unit 0 description R3-to-R6
set interfaces ge-1/2/20 unit 0 family inet address 1.2.96.2/30
set interfaces ge-1/2/20 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.3/32 primary
set routing-options autonomous-system 10
set routing-options multicast stream-protection
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.3
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.1

```



```

set protocols bgp group ibgp neighbor 1.1.1.5
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/18.0
set protocols pim interface ge-1/2/22.0
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
  orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 1.1.0.1/30 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept

```

Device R4

```

set interfaces ge-1/2/18 unit 0 description R4-to-R3
set interfaces ge-1/2/18 unit 0 family inet address 1.3.4.2/30
set interfaces ge-1/2/18 unit 0 family mpls
set interfaces ge-1/2/23 unit 0 description R4-to-R7
set interfaces ge-1/2/23 unit 0 family inet address 1.4.7.1/30
set interfaces lo0 unit 0 family inet address 1.1.1.4/32
set protocols igmp interface ge-1/2/18.0 version 3
set protocols igmp interface ge-1/2/18.0 static group 232.1.1.1 group-count 2
set protocols igmp interface ge-1/2/18.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface ge-1/2/18.0 static group 232.2.2.2 source 1.2.7.7
set protocols sap listen 232.1.1.1
set protocols sap listen 232.2.2.2
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all

```



```

set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface ge-1/2/23.0
set protocols pim interface ge-1/2/18.0
set protocols pim interface lo0.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
  orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 10

```

Device R5

```

set interfaces ge-1/2/24 unit 0 description R5-to-src2
set interfaces ge-1/2/24 unit 0 family inet address 1.5.0.1/30
set interfaces ge-1/2/16 unit 0 description R5-to-R2
set interfaces ge-1/2/16 unit 0 family inet address 1.2.5.2/30
set interfaces ge-1/2/16 unit 0 family mpls
set interfaces ge-1/2/25 unit 0 description R5-to-R6
set interfaces ge-1/2/25 unit 0 family inet address 1.5.6.1/30
set interfaces ge-1/2/25 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.5/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.5
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.7
set protocols bgp group ibgp neighbor 1.1.1.3
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/16.0
set protocols ldp interface ge-1/2/25.0
set protocols ldp p2mp

```



```

set protocols pim interface lo0.0
set protocols pim interface ge-1/2/24.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 10

```

Device R6

```

set interfaces ge-1/2/13 unit 0 description R6-to-R1
set interfaces ge-1/2/13 unit 0 family inet address 1.1.6.2/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/19 unit 0 description R6-to-R3
set interfaces ge-1/2/19 unit 0 family inet address 1.3.6.1/30
set interfaces ge-1/2/19 unit 0 family mpls
set interfaces ge-1/2/25 unit 0 description R6-to-R5
set interfaces ge-1/2/25 unit 0 family inet address 1.5.6.2/30
set interfaces ge-1/2/25 unit 0 family mpls
set interfaces ge-1/2/26 unit 0 description R6-to-R7
set interfaces ge-1/2/26 unit 0 family inet address 1.6.7.1/30
set interfaces ge-1/2/26 unit 0 family mpls
set interfaces ge-1/2/20 unit 0 description R6-to-R3
set interfaces ge-1/2/20 unit 0 family inet address 1.2.96.1/30
set interfaces ge-1/2/20 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.6/30
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp

```

Device R7

```

set interfaces ge-1/2/17 unit 0 description R7-to-R2
set interfaces ge-1/2/17 unit 0 family inet address 1.2.7.2/30

```



```

set interfaces ge-1/2/17 unit 0 family mpls
set interfaces ge-1/2/21 unit 0 description R7-to-R3
set interfaces ge-1/2/21 unit 0 family inet address 1.3.7.2/30
set interfaces ge-1/2/21 unit 0 family mpls
set interfaces ge-1/2/23 unit 0 description R7-to-R4
set interfaces ge-1/2/23 unit 0 family inet address 1.4.7.2/30
set interfaces ge-1/2/23 unit 0 family mpls
set interfaces ge-1/2/26 unit 0 description R7-to-R6
set interfaces ge-1/2/26 unit 0 family inet address 1.6.7.2/30
set interfaces ge-1/2/26 unit 0 family mpls
set interfaces ge-1/2/27 unit 0 description R7-to-R8
set interfaces ge-1/2/27 unit 0 family inet address 1.7.8.1/30
set interfaces ge-1/2/27 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.7/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.7
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols bgp group ibgp neighbor 1.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/17.0
set protocols ldp interface ge-1/2/21.0
set protocols ldp interface ge-1/2/26.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/27.0
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
    orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
    orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 1.1.0.1/30 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 10

```



```
set routing-options multicast stream-protection policy mldppim-ex
```

Device R8

```
set interfaces ge-1/2/22 unit 0 description R8-to-R3
set interfaces ge-1/2/22 unit 0 family inet address 1.3.8.2/30
set interfaces ge-1/2/22 unit 0 family mpls
set interfaces ge-1/2/27 unit 0 description R8-to-R7
set interfaces ge-1/2/27 unit 0 family inet address 1.7.8.2/30
set interfaces ge-1/2/27 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.8/32
set protocols igmp interface ge-1/2/22.0 version 3
set protocols igmp interface ge-1/2/22.0 static group 232.1.1.1 group-count 2
set protocols igmp interface ge-1/2/22.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface ge-1/2/22.0 static group 232.2.2.2 source 1.2.7.7
set protocols sap listen 232.1.1.1
set protocols sap listen 232.2.2.2
set protocols rsvp interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface ge-1/2/27.0
set protocols pim interface ge-1/2/22.0
set protocols pim interface lo0.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
    orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
    orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 10
```

Configuration

Step-by-Step Procedure

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

To configure Device R3:

1. Enable enhanced IP mode.

```
[edit chassis]
user@R3# set network-services enhanced-ip
```

2. Configure the device interfaces.

```
[edit interfaces]
user@R3# set ge-1/2/14 unit 0 description R3-to-R2
user@R3# set ge-1/2/14 unit 0 family inet address 1.2.3.2/30
user@R3# set ge-1/2/14 unit 0 family mpls
user@R3# set ge-1/2/18 unit 0 description R3-to-R4
user@R3# set ge-1/2/18 unit 0 family inet address 1.3.4.1/30
user@R3# set ge-1/2/18 unit 0 family mpls
user@R3# set ge-1/2/19 unit 0 description R3-to-R6
user@R3# set ge-1/2/19 unit 0 family inet address 1.3.6.2/30
user@R3# set ge-1/2/19 unit 0 family mpls
user@R3# set ge-1/2/21 unit 0 description R3-to-R7
user@R3# set ge-1/2/21 unit 0 family inet address 1.3.7.1/30
user@R3# set ge-1/2/21 unit 0 family mpls
user@R3# set ge-1/2/22 unit 0 description R3-to-R8
user@R3# set ge-1/2/22 unit 0 family inet address 1.3.8.1/30
user@R3# set ge-1/2/22 unit 0 family mpls
user@R3# set ge-1/2/15 unit 0 description R3-to-R2
user@R3# set ge-1/2/15 unit 0 family inet address 1.2.94.2/30
user@R3# set ge-1/2/15 unit 0 family mpls
user@R3# set ge-1/2/20 unit 0 description R3-to-R6
user@R3# set ge-1/2/20 unit 0 family inet address 1.2.96.2/30
user@R3# set ge-1/2/20 unit 0 family mpls
user@R3# set lo0 unit 0 family inet address 1.1.1.3/32 primary
```

3. Configure the autonomous system (AS) number.

```
user@R3# set routing-options autonomous-system 10
```

4. Configure the routing policies.


```
[edit policy-options policy-statement mldppim-ex]
user@R3# set term B from source-address-filter 192.168.0.0/24 orlonger
user@R3# set term B from source-address-filter 192.168.219.11/32 orlonger
user@R3# set term B then accept
user@R3# set term A from source-address-filter 1.1.0.1/30 orlonger
user@R3# set term A then accept
[edit policy-options policy-statement static-route-tobgp]
user@R3# set term static from protocol static
user@R3# set term static from protocol direct
user@R3# set term static then accept
```

5. Configure PIM.

```
[edit protocols pim]
user@R3# set mldp-inband-signalling policy mldppim-ex
user@R3# set interface lo0.0
user@R3# set interface ge-1/2/18.0
user@R3# set interface ge-1/2/22.0
```

6. Configure LDP.

```
[edit protocols ldp]
user@R3# set interface all
user@R3# set p2mp
```

7. Configure an IGP or static routes.

```
[edit protocols ospf]
user@R3# set traffic-engineering
user@R3# set area 0.0.0.0 interface all
user@R3# set area 0.0.0.0 interface fxp0.0 disable
user@R3# set area 0.0.0.0 interface lo0.0 passive
```

8. Configure internal BGP.

```
[edit protocols bgp group ibgp]
user@R3# set local-address 1.1.1.3
user@R3# set peer-as 10
user@R3# set neighbor 1.1.1.1
user@R3# set neighbor 1.1.1.5
```


9. Configure MPLS and, optionally, RSVP.

```
[edit protocols mpls]
user@R3# set interface all
[edit protocols rsvp]
user@R3# set interface all
```

10. Enable MoFRR.

```
[edit routing-options multicast]
user@R3# set stream-protection
```

Results

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

```
user@R3# show chassis
network-services enhanced-ip;
```

```
user@R3# show interfaces
ge-1/2/14 {
  unit 0 {
    description R3-to-R2;
    family inet {
      address 1.2.3.2/30;
    }
    family mpls;
  }
}
ge-1/2/18 {
  unit 0 {
    description R3-to-R4;
    family inet {
      address 1.3.4.1/30;
    }
    family mpls;
  }
}
ge-1/2/19 {
  unit 0 {
```



```

        description R3-to-R6;
        family inet {
            address 1.3.6.2/30;
        }
        family mpls;
    }
}
ge-1/2/21 {
    unit 0 {
        description R3-to-R7;
        family inet {
            address 1.3.7.1/30;
        }
        family mpls;
    }
}
ge-1/2/22 {
    unit 0 {
        description R3-to-R8;
        family inet {
            address 1.3.8.1/30;
        }
        family mpls;
    }
}
ge-1/2/15 {
    unit 0 {
        description R3-to-R2;
        family inet {
            address 1.2.94.2/30;
        }
        family mpls;
    }
}
ge-1/2/20 {
    unit 0 {
        description R3-to-R6;
        family inet {
            address 1.2.96.2/30;
        }
        family mpls;
    }
}
lo0 {

```



```

unit 0 {
    family inet {
        address 192.168.15.1/32;
        address 1.1.1.3/32 {
            primary;
        }
    }
}

```

user@R3# **show protocols**

```

rsvp {
    interface all;
}
mpls {
    interface all;
}
bgp {
    group ibgp {
        local-address 1.1.1.3;
        peer-as 10;
        neighbor 1.1.1.1;
        neighbor 1.1.1.5;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface all;
    p2mp;
}
pim {
    mldp-inband-signalling {
        policy mldppim-ex;
    }
}

```



```

}
interface lo0.0;
interface ge-1/2/18.0;
interface ge-1/2/22.0;
}

```

```

user@R3# show policy-options
policy-statement mldppim-ex {
  term B {
    from {
      source-address-filter 192.168.0.0/24 orlonger;
      source-address-filter 192.168.219.11/32 orlonger;
    }
    then accept;
  }
  term A {
    from {
      source-address-filter 1.1.0.1/30 orlonger;
    }
    then accept;
  }
}
policy-statement static-route-tobgp {
  term static {
    from protocol [ static direct ];
    then accept;
  }
}

```

```

user@R3# show routing-options
autonomous-system 10;
multicast {
  stream-protection;
}

```

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

Verification

IN THIS SECTION

- [Checking the LDP Point-to-Multipoint Forwarding Equivalency Classes | 1041](#)
- [Examining the Label Information | 1041](#)
- [Checking the Multicast Routes | 1043](#)
- [Checking the LDP Point-to-Multipoint Traffic Statistics | 1045](#)

Confirm that the configuration is working properly.

Checking the LDP Point-to-Multipoint Forwarding Equivalency Classes

Purpose

Make sure the MoFRR is enabled, and determine what labels are being used.

Action

```
user@R3> show ldp p2mp fec
```

```
LDP P2MP FECs:
P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
MoFRR enabled
Fec type: Egress (Active)
Label: 301568
P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
MoFRR enabled
Fec type: Egress (Active)
Label: 301600
```

Meaning

The output shows that MoFRR is enabled, and it shows that the labels 301568 and 301600 are being used for the two multipoint LDP point-to-multipoint LSPs.

Examining the Label Information

Purpose

Make sure that the egress device has two upstream interfaces for the multicast group join.

Action

user@R3> **show route label 301568 detail**

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Address: 0x2735208
              Next-hop reference count: 3
              Next hop type: Router, Next hop index: 1397
              Address: 0x2735d2c
              Next-hop reference count: 3
              Next hop: 1.3.8.2 via ge-1/2/22.0
              Label operation: Pop
              Load balance label: None;
              Next hop type: Router, Next hop index: 1395
              Address: 0x2736290
              Next-hop reference count: 3
              Next hop: 1.3.4.2 via ge-1/2/18.0
              Label operation: Pop
              Load balance label: None;
              State: <Active Int AckRequest MulticastRPF>
              Local AS:      10
              Age: 54:05      Metric: 1
              Validation State: unverified
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
Primary Upstream : 1.1.1.3:0--1.1.1.2:0
    RPF Nexthops :
        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
Backup Upstream : 1.1.1.3:0--1.1.1.6:0
    RPF Nexthops :
        ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
        ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe

```

user@R3> **show route label 301600 detail**


```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301600 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Address: 0x27356b4
              Next-hop reference count: 3
              Next hop type: Router, Next hop index: 1520
              Address: 0x27350f4
              Next-hop reference count: 3
              Next hop: 1.3.8.2 via ge-1/2/22.0
              Label operation: Pop
              Load balance label: None;
              Next hop type: Router, Next hop index: 1481
              Address: 0x273645c
              Next-hop reference count: 3
              Next hop: 1.3.4.2 via ge-1/2/18.0
              Label operation: Pop
              Load balance label: None;
              State: <Active Int AckRequest MulticastRPF>
              Local AS:      10
              Age: 54:25      Metric: 1
              Validation State: unverified
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src:
192.168.219.11
Primary Upstream : 1.1.1.3:0--1.1.1.6:0
    RPF Nexthops :
        ge-1/2/20.0, 1.2.96.1, Label: 301600, weight: 0x1
        ge-1/2/19.0, 1.3.6.1, Label: 301600, weight: 0x1
Backup Upstream : 1.1.1.3:0--1.1.1.2:0
    RPF Nexthops :
        ge-1/2/15.0, 1.2.94.1, Label: 301616, weight: 0xfffe
        ge-1/2/14.0, 1.2.3.1, Label: 301616, weight: 0xfffe

```

Meaning

The output shows the primary upstream paths and the backup upstream paths. It also shows the RPF next hops.

Checking the Multicast Routes

Purpose

Examine the IP multicast forwarding table to make sure that there is an upstream RPF interface list, with a primary and a backup interface.

Action

user@R3> **show ldp p2mp path**

```
P2MP path type: Transit/Egress
  Output Session (label): 1.1.1.2:0 (301568) (Primary)
  Egress Nexthops: Interface ge-1/2/18.0
                   Interface ge-1/2/22.0
  RPF Nexthops:   Interface ge-1/2/15.0, 1.2.94.1, 301568, 1
                   Interface ge-1/2/20.0, 1.2.96.1, 301584, 65534
                   Interface ge-1/2/14.0, 1.2.3.1, 301568, 1
                   Interface ge-1/2/19.0, 1.3.6.1, 301584, 65534
  Attached FECs:  P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
(Active)
P2MP path type: Transit/Egress
  Output Session (label): 1.1.1.6:0 (301584) (Backup)
  Egress Nexthops: Interface ge-1/2/18.0
                   Interface ge-1/2/22.0
  RPF Nexthops:   Interface ge-1/2/15.0, 1.2.94.1, 301568, 1
                   Interface ge-1/2/20.0, 1.2.96.1, 301584, 65534
                   Interface ge-1/2/14.0, 1.2.3.1, 301568, 1
                   Interface ge-1/2/19.0, 1.3.6.1, 301584, 65534
  Attached FECs:  P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
(Active)
P2MP path type: Transit/Egress
  Output Session (label): 1.1.1.6:0 (301600) (Primary)
  Egress Nexthops: Interface ge-1/2/18.0
                   Interface ge-1/2/22.0
  RPF Nexthops:   Interface ge-1/2/15.0, 1.2.94.1, 301616, 65534
                   Interface ge-1/2/20.0, 1.2.96.1, 301600, 1
                   Interface ge-1/2/14.0, 1.2.3.1, 301616, 65534
                   Interface ge-1/2/19.0, 1.3.6.1, 301600, 1
  Attached FECs:  P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
(Active)
P2MP path type: Transit/Egress
  Output Session (label): 1.1.1.2:0 (301616) (Backup)
  Egress Nexthops: Interface ge-1/2/18.0
                   Interface ge-1/2/22.0
  RPF Nexthops:   Interface ge-1/2/15.0, 1.2.94.1, 301616, 65534
                   Interface ge-1/2/20.0, 1.2.96.1, 301600, 1
                   Interface ge-1/2/14.0, 1.2.3.1, 301616, 65534
                   Interface ge-1/2/19.0, 1.3.6.1, 301600, 1
```



```
Attached FECs: P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
(Active)
```

Meaning

The output shows primary and backup sessions, and RPF next hops.

Checking the LDP Point-to-Multipoint Traffic Statistics

Purpose

Make sure that both primary and backup statistics are listed.

Action

```
user@R3> show ldp traffic-statistics p2mp
```

```
P2MP FEC Statistics:
```

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

Meaning

The output shows both primary and backup routes with the labels.

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

[Configuring Multicast-Only Fast Reroute | 1021](#)

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring LDP Downstream on Demand

IN THIS SECTION

- [Requirements | 1046](#)
- [Overview | 1047](#)
- [Configuration | 1047](#)
- [Verification | 1051](#)

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

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_ldap;
}

```

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 LDP Native IPv6 Support

LDP is supported in an IPv6-only network, and in an IPv6 or IPv4 dual-stack network as described in *RFC 7552*. Configure the address family as **inet** for IPv4 or **inet6** for IPv6 or both, and the transport preference to be either **IPv4** or **IPv6**. The **dual-transport** statement allows Junos OS LDP to establish the TCP connection over IPv4 with IPv4 neighbors, and over IPv6 with IPv6 neighbors as a single-stack LSR. The **inet-lsr-id** and **inet6-lsr-id** IDs are the two LSR IDs that have to be configured to establish an LDP session over IPv4 and IPv6 TCP transport. These two IDs should be non-zero and must be configured with different values.

Before you configure IPv6 as dual-stack, be sure you configure the routing and signaling protocols.

To configure LDP native IPv6 support, you must do the following:

1. Enable forwarding equivalence class (FEC) deaggregation in order to use different labels for different address families.

```
[edit protocols ldp]  
set deaggregate
```

2. Configure LDP address families.

```
[edit protocols ldp]  
set family inet6  
set family inet
```

3. Configure the **transport-preference** statement to select the preferred transport for the TCP connection when both IPv4 and IPv6 are enabled. By default, IPv6 is used as the TCP transport for establishing an LDP connection.

```
[edit protocols ldp]  
set transport-preference ipv4
```

4. (Optional) Configure dual-transport to allow LDP to establish a separate IPv4 session with an IPv4 neighbor, and an IPv6 session with an IPv6 neighbor. Configure **inet-lsr-id** as the LSR ID for IPv4, and **inet6-lsr-id** as the LSR ID for IPv6.

```
[edit protocols ldp dual-transport]  
set inet-lsr-id inet-lsr-id  
set inet6-lsr-id inet6-lsr-id
```


For example, configure `inet-lsr-id` as 10.255.0.1, and `inet6-lsr-id` as 1.1.1.1.

```
[edit protocols ldp dual-transport]
set inet-lsr-id 10.255.0.1
set inet6-lsr-id 1.1.1.1
```

RELATED DOCUMENTATION

[LDP Native IPv6 Support Overview | 932](#)

[Example: Configuring LDP Native IPv6 Support | 1054](#)

Example: Configuring LDP Native IPv6 Support

IN THIS SECTION

- [Requirements | 1054](#)
- [Overview | 1055](#)
- [Configuration | 1055](#)

This example shows how to allow the Junos OS Label Distribution Protocol (LDP) to establish the TCP connection over IPv4 with IPv4 neighbors, and over IPv6 with IPv6 neighbors as a single-stack LSR. This helps avoid tunneling of IPv6 over IPv4 MPLS core with IPv4-signaled MPLS label-switched paths (LSPs).

Requirements

This example uses the following hardware and software components:

- Two MX Series routers
- Junos OS Release 16.1 or later running on all devices

Before you configure IPv6 as dual-stack, be sure you configure the routing and signaling protocols.

Overview

LDP is supported in an IPv6 only network, and in an IPv6 or IPv4 dual-stack network as described in *RFC 7552*. Configure the address family as **inet** for IPv4 or **inet6** for IPv6. By default, IPv6 is used as the TCP transport for the LDP session with its peers when both IPv4 and IPv6 are enabled. The dual-transport statement allows Junos LDP to establish the TCP connection over IPv4 with IPv4 neighbors, and over IPv6 with IPv6 neighbors as a single-stack LSR. The **inet-lsr-id** and **inet6-lsr-id** are the two LSR IDs that have to be configured to establish an LDP session over IPv4 and IPv6 TCP transport. These two IDs should be non-zero and must be configured with different values.

Topology

Figure 88 on page 1055 shows the LDP IPv6 configured as dual-stack on Device R1 and Device R2.

Figure 88: Example LDP Native IPv6 Support



Configuration

CLI Quick Configuration

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

R1

```

set interfaces ge-1/0/0 unit 0 family inet address 192.168.12.1/24
set interfaces ge-1/0/0 unit 0 family iso
set interfaces ge-1/0/0 unit 0 family inet6 address 2001:db8:0:12::/64 eui-64
set interfaces ge-1/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.1/32
set interfaces lo0 unit 0 family iso address 49.0001.1720.1600.1010.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::1/128
set protocols isis interface ge-1/0/0.0
set protocols isis interface lo0.0
  
```



```

set protocols mpls interface ge-1/0/0.0
set protocols ldp deaggregate
set protocols ldp interface ge-1/0/0.0
set protocols ldp interface lo0.0
set protocols ldp family inet6
set protocols ldp family inet

```

R2

```

set interfaces ge-1/0/1 unit 0 family inet address 192.168.12.2/24
set interfaces ge-1/0/1 unit 0 family iso
set interfaces ge-1/0/1 unit 0 family inet6 address 2001:db8:0:12::/64 eui-64
set interfaces ge-1/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.2/32
set interfaces lo0 unit 0 family iso address 49.0001.1720.1600.2020.00
set interfaces lo0 unit 0 family inet6 address 2001:db8::2/128
set protocols isis interface ge-1/0/1.0
set protocols isis interface lo0.0
set protocols mpls interface ge-1/0/1.0
set protocols ldp deaggregate
set protocols ldp interface ge-1/0/1.0
set protocols ldp interface lo0.0
set protocols ldp family inet6
set protocols ldp family inet

```

Configuring R1

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see “Using the CLI Editor in Configuration Mode” in the *CLI User Guide*.

To configure Device R1:

1. Configure the interfaces.

```

[edit interfaces]
set ge-1/0/0 unit 0 family inet address 192.168.12.1/24
set ge-1/0/0 unit 0 family iso

```



```
set ge-1/0/0 unit 0 family inet6 address 2001:db8:0:12::/64 eui-64
set ge-1/0/0 unit 0 family mpls
```

2. Assign a loopback address to the device.

```
[edit interfaces lo0 unit 0]
set family inet address 10.255.0.1/32
set family iso address 49.0001.1720.1600.1010.00
set family inet6 address 2001:db8::1/128
```

3. Configure the IS-IS interfaces.

```
[edit protocols isis]
set interface ge-1/0/0.0
set interface lo0.0
```

4. Configure MPLS to use LDP interfaces on the device.

```
[edit protocols mpls]
set protocols mpls interface ge-1/0/0.0
set interface ge-1/0/0.0
set interface lo0.0
```

5. Enable forwarding equivalence class (FEC) deaggregation in order to use different labels for different address families.

```
[edit protocols ldp]
set deaggregate
```

6. Configure LDP address families.

```
[edit protocols ldp]
set family inet6
set family inet
```

Results

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

```
user@R1# show interfaces
ge-1/0/0 {
  unit 0 {
    family inet {
      address 192.168.12.1/24;
    }
    family iso;
    family inet6 {
      address 2001:db8:0:12::/64 {
        eui-64;
      }
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.1/32;
    }
    family iso {
      address 49.0001.1720.1600.1010.00
    }
    family inet6 {
      address 2001:db8::1/128;
    }
  }
}
```

```
user@R1# show protocols
mpls {
  interface ge-1/0/0.0;
}
isis {
  interface ge-1/0/0.0;
  interface lo0.0;
}
ldp {
```



```

deaggregate;
interface ge-1/0/0.0;
interface lo0.0;
family {
    inet6;
    inet;
}
}

```

Verification

IN THIS SECTION

- [Verifying the Route Entries in the mpls.0 Table | 1059](#)
- [Verifying the Route Entries in the inet.3 Table | 1060](#)
- [Verifying the Route Entries in the inet6.3 Table | 1060](#)
- [Verifying the LDP Database | 1061](#)
- [Verifying the LDP Neighbor Information | 1062](#)
- [Verifying the LDP Session Information | 1063](#)

Confirm that the configuration is working properly.

Verifying the Route Entries in the mpls.0 Table

Purpose

Display mpls.0 route table information.

Action

On Device R1, from operational mode, run the **show route table mpls.0** command to display mpls.0 route table information.

user@R1> **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] 05:19:58, metric 1
            Receive
1          *[MPLS/0] 05:19:58, metric 1

```



```

                Receive
2                *[MPLS/0] 05:19:58, metric 1
                Receive
13              *[MPLS/0] 05:19:58, metric 1
                Receive
299824          *[LDP/9] 04:28:45, metric 1
                > to fe80::21f:1200:cb6:4c8d via ge-1/0/0.0, Pop
299824(S=0)     *[LDP/9] 04:28:45, metric 1
                > to fe80::21f:1200:cb6:4c8d via ge-1/0/0.0, Pop
299888          *[LDP/9] 00:56:12, metric 1
                > to 192.168.12.2 via ge-1/0/0.0, Pop
299888(S=0)     *[LDP/9] 00:56:12, metric 1
                > to 192.168.12.2 via ge-1/0/0.0, Pop

```

Meaning

The output shows the mpls.0 route table information.

Verifying the Route Entries in the inet.3 Table

Purpose

Display inet.3 route table information.

Action

On Device R1, from operational mode, run the **show route table inet.3** command to display inet.3 route table information.

```
user@R1> show route table inet.3
```

```

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.0.2/32      *[LDP/9] 00:58:38, metric 1
                   > to 192.168.12.2 via ge-1/0/0.0

```

Meaning

The output shows the inet.3 route table information.

Verifying the Route Entries in the inet6.3 Table

Purpose

Display inet6.3 route table information.

Action

On Device R1, from operational mode, run the **show route table inet6.3** command to display inet6.3 route table information.

```
user@R1> show route table inet6.3
```

```
inet6.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::2/128      *[LDP/9] 04:31:17, metric 1
                    > to fe80::21f:1200:cb6:4c8d via ge-1/0/0.0
```

Meaning

The output shows the inet6.3 route table information.

Verifying the LDP Database

Purpose

Display the LDP database information.

Action

On Device R1, from operational mode, run the **show ldp database** command to display LDP database information.

```
user@R1> show ldp database
```

```
Input label database, 10.255.0.1:0--10.255.0.2:0
Labels received: 3
  Label      Prefix
  299840     10.255.0.1/32
    3        10.255.0.2/32
  299808     2001:db8::1/128
    3        2001:db8::2/128

Output label database, 10.255.0.1:0--10.255.0.2:0
Labels advertised: 3
  Label      Prefix
    3        10.255.0.1/32
  299888     10.255.0.2/32
```


3	2001:db8::1/128
299824	2001:db8::2/128

Meaning

The output shows the entries in the LDP database.

Verifying the LDP Neighbor Information

Purpose

Display the LDP neighbor information.

Action

On Device R1, from operational mode, run the **show ldp neighbor** and **show ldp neighbor extensive** commands to display LDP neighbor information.

```
user@R1> show ldp neighbor
```

Address	Interface	Label space ID	Hold time
fe80::21f:1200:cb6:4c8d	ge-1/0/0.0	10.255.0.2:0	12
192.168.12.2	ge-1/0/0.0	10.255.0.2:0	11

```
user@R1> show ldp neighbor extensive
```

Address	Interface	Label space ID	Hold time
192.168.12.2	ge-1/0/0.0	10.255.0.2:0	11
Transport address: 10.255.0.2, Transport preference: IPv6, Configuration sequence: 10 Up for 00:04:35 Reference count: 1 Hold time: 15, Proposed local/peer: 15/15 Hello flags: none Neighbor types: discovered			
Address	Interface	Label space ID	Hold time
fe80::21f:1200:cb6:4c8d	ge-1/0/0.0	10.255.0.2:0	14
Transport address: 2001:db8::2, Transport preference: IPv6, Configuration sequence: 10 Up for 00:04:35 Reference count: 1 Hold time: 15, Proposed local/peer: 15/15 Hello flags: none			


```
Neighbor types: discovered
```

Meaning

The output shows LDP neighbor information of both IPv4 and IPv6 addresses.

Verifying the LDP Session Information

Purpose

Display the LDP session information.

Action

On Device R1, from operational mode, run the **show ldp session** and **show ldp session extensive** commands to display LDP session information.

```
user@R1> show ldp session
```

```
session
  Address                               State      Connection Hold time  Adv. Mode
2001:db8::2                           Operational Open        20         DU
```

```
user@R1> show ldp session extensive
```

```
Address: 2001:db8::2, State: Operational, Connection: Open, Hold time: 29
Session ID: 10.255.0.1:0--10.255.0.2:0
Next keepalive in 9 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 2001:db8::1, Remote address: 2001:db8::2
Up for 00:05:31
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Session flags: none
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
MTU discovery: disabled
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.255.0.2
  192.168.12.2
  2001:db8::2
  fe80::21f:1200:cb6:4c8d
Queue depth: 0

```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Initialization	1	1	0	0
Keepalive	34	34	0	0
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	3	3	0	0
Label request	0	0	0	0
Label withdraw	0	0	0	0
Label release	0	0	0	0
Label abort	0	0	0	0

Meaning

The output displays information for the LDP session using IPv6 as the TCP transport.

Configure transport-preference to Select the Preferred Transport

CLI Quick Configuration

You can configure the **transport-preference** statement to select the preferred transport for a TCP connection when both IPv4 and IPv6 are enabled. By default, IPv6 is used as TCP transport for establishing an LDP connection.

- (Optional) Configure the transport preference for an LDP connection.

```

[edit protocols ldp]
set transport-preference ipv4

```

Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
mpls {
  interface ge-1/0/0.0;
}
isis {
  interface ge-1/0/0.0;
  interface lo0.0;
}
ldp {
  deaggregate;
  interface ge-1/0/0.0;
  interface lo0.0;
  family {
    inet6;
    inet;
  }
  transport-preference ipv4;
}
```

Verification

IN THIS SECTION

- [Verifying the LDP Neighbor Information | 1065](#)
- [Verifying the LDP Session Information | 1066](#)

Confirm that the configuration is working properly.

Verifying the LDP Neighbor Information

Purpose

Display the LDP neighbor information.

Action

On Device R1, from operational mode, run the **show ldp neighbor extensive** command to display LDP neighbor information.

```
user@R1> show ldp neighbor extensive
```

```

Address                Interface      Label space ID      Hold time
192.168.12.2           ge-1/0/0.0    10.255.0.2:0        14
  Transport address: 10.255.0.2, Transport preference: IPv4, Configuration sequence:
  9
  Up for 00:00:14
  Reference count: 1
  Hold time: 15, Proposed local/peer: 15/15
  Hello flags: none
  Neighbor types: discovered
Address                Interface      Label space ID      Hold time
fe80::21f:1200:cb6:4c8d ge-1/0/0.0    10.255.0.2:0        14
  Transport address: 2001:db8::2, Transport preference: IPv4, Configuration
sequence: 9
  Up for 00:00:14
  Reference count: 1
  Hold time: 15, Proposed local/peer: 15/15
  Hello flags: none
  Neighbor types: discovered
```

Meaning

The output shows LDP neighbor information for both the IPv4 and IPv6 addresses.

Verifying the LDP Session Information

Purpose

Display the LDP session information.

Action

On Device R1, from operational mode, run the **show ldp session extensive** command to display LDP session information.

```
user@R1> show ldp session extensive
```

```

Address: 10.255.0.2, State: Operational, Connection: Open, Hold time: 24
  Session ID: 10.255.0.1:0--10.255.0.2:0
  Next keepalive in 4 seconds
```



```

Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 2
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 10.255.0.1, Remote address: 10.255.0.2
Up for 00:05:26
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Session flags: none
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
MTU discovery: disabled
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.255.0.2
  192.168.12.2
  2001:db8::2
  fe80::21f:1200:cb6:4c8d
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	2	2	0	0
Address withdraw	0	0	0	0
Label mapping	6	6	0	0
Label request	0	0	0	0
Label withdraw	0	0	0	0
Label release	0	0	0	0
Label abort	0	0	0	0

Meaning

The output displays information for the LDP session using IPv6 as the TCP transport.

Configure dual-transport to Establish Separate Sessions for IPv4 with an IPv4 Neighbor and IPv6 with an IPv6 Neighbor

Step-by-Step Procedure

You can configure the **dual-transport** statement to allow LDP to establish a separate IPv4 session with an IPv4 neighbor, and an IPv6 session with an IPv6 neighbor. This requires the configuration of **inet-lsr-id** as the LSR ID for IPv4, and **inet6-lsr-id** as the LSR ID for IPv6.

- (Optional) Configure dual-transport to allow LDP to establish the TCP connection over IPv4 with IPv4 neighbors, and over IPv6 with IPv6 neighbors as a single-stack LSR.

```
[edit protocols ldp dual-transport]
set inet-lsr-id 10.255.0.1
set inet6-lsr-id 1.1.1.1
```

Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show protocols
mpls {
  interface ge-1/0/0.0;
}
isis {
  interface ge-1/0/0.0;
  interface lo0.0;
}
ldp {
  deaggregate;
  interface ge-1/0/0.0;
  interface lo0.0;
  family {
    inet6;
    inet;
  }
  dual-transport {
    inet-lsr-id 10.255.0.1;
    inet6-lsr-id 1.1.1.1;
  }
}
```


Verification

IN THIS SECTION

- [Verifying the LDP Neighbor Information | 1069](#)
- [Verifying the LDP Session Information | 1070](#)

Confirm that the configuration is working properly.

Verifying the LDP Neighbor Information

Purpose

Display the LDP neighbor information.

Action

On Device R1, from operational mode, run the **show ldp neighbor extensive** command to display LDP neighbor information.

```
user@R1> show ldp neighbor extensive
```

```

Address                Interface      Label space ID      Hold time
192.168.12.2           ge-1/0/0.0    10.255.0.2:0        11
  Transport address: 10.255.0.2, Configuration sequence: 10
  Up for 00:04:35
  Reference count: 1
  Hold time: 15, Proposed local/peer: 15/15
  Hello flags: none
  Neighbor types: discovered
Address                Interface      Label space ID      Hold time
fe80::21f:1200:cb6:4c8d ge-1/0/0.0    10.255.0.2:0        14
  Transport address: 2001:db8::2, Configuration sequence: 10
  Up for 00:04:35
  Reference count: 1
  Hold time: 15, Proposed local/peer: 15/15
  Hello flags: none
  Neighbor types: discovered
```

Meaning

The output shows LDP neighbor information for both the IPv4 and IPv6 addresses.

Verifying the LDP Session Information

Purpose

Display the LDP session information.

Action

On Device R1, from operational mode, run the **show ldp session extensive** command to display LDP neighbor information.

user@R1> **show ldp session extensive**

```
Address: 2001:db8::2, State: Operational, Connection: Open, Hold time: 29
  Session ID: 1.1.1.1:0--10.255.0.2:0
  Next keepalive in 9 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 2001:db8::1, Remote address: 2001:db8::2
  Up for 00:05:31
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Session flags: none
  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
  MTU discovery: disabled
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    2001:db8::2
    fe80::21f:1200:cb6:4c8d
  Queue depth: 0
```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Initialization	1	1	0	0
Keepalive	34	34	0	0
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	3	3	0	0

Label request	0	0	0	0
Label withdraw	0	0	0	0
Label release	0	0	0	0
Label abort	0	0	0	0

Address: 10.255.0.2, State: Operational, Connection: Open, Hold time: 29

Session ID: 10.255.0.1:0--10.255.0.2:0

Next keepalive in 9 seconds

Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1

Neighbor types: discovered

Keepalive interval: 10, Connect retry interval: 1

Local address: 10.255.0.1, Remote address: 10.255.0.2

Up for 00:05:31

Capabilities advertised: none

Capabilities received: none

Protection: disabled

Session flags: none

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

MTU discovery: disabled

Nonstop routing state: Not in sync

Next-hop addresses received:

10.255.0.2

192.168.12.2

Queue depth: 0

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Initialization	1	1	0	0
Keepalive	34	34	0	0
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	3	3	0	0
Label request	0	0	0	0
Label withdraw	0	0	0	0
Label release	0	0	0	0
Label abort	0	0	0	0

RELATED DOCUMENTATION

[LDP Native IPv6 Support Overview | 932](#)

[Configuring LDP Native IPv6 Support | 1053](#)

[family \(Protocols LDP\) | 2614](#)

[dual-transport | 2592](#)

Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs

IN THIS SECTION

- [Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs | 1072](#)
- [Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs | 1082](#)

Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs

IN THIS SECTION

- [How M-LDP Works | 1073](#)
- [Terminology | 1078](#)
- [Ingress Join Translation and Pseudo Interface Handling | 1079](#)
- [Ingress Splicing | 1079](#)
- [Reverse Path Forwarding | 1079](#)
- [LSP Root Detection | 1080](#)
- [Egress Join Translation and Pseudo Interface Handling | 1080](#)
- [Egress Splicing | 1080](#)
- [Supported Functionality | 1080](#)
- [Unsupported Functionality | 1081](#)
- [LDP Functionality | 1081](#)
- [Egress LER Functionality | 1081](#)
- [Transit LSR Functionality | 1082](#)
- [Ingress LER Functionality | 1082](#)

The Multipoint Label Distribution Protocol (M-LDP) for point-to-multipoint label-switched paths (LSPs) with in-band signaling is useful in a deployment with an existing IP/MPLS backbone, in which you need to carry multicast traffic, for IPTV for example.

For years, the most widely used solution for transporting multicast traffic has been to use native IP multicast in the service provider core with multipoint IP tunneling to isolate customer traffic. A multicast routing protocol, usually Protocol Independent Multicast (PIM), is deployed to set up the forwarding paths. IP multicast routing is used for forwarding, using PIM signaling in the core. For this model to work, the core network has to be multicast enabled. This allows for effective and stable deployments even in inter-autonomous system (AS) scenarios.

However, in an existing IP/MPLS network, deploying PIM might not be the first choice. Some service providers are interested in replacing IP tunneling with MPLS label encapsulation. The motivations for moving to MPLS label switching is to leverage MPLS traffic engineering and protection features and to reduce the amount of control traffic overhead in the provider core.

To do this, service providers are interested in leveraging the extension of the existing deployments to allow multicast traffic to pass through. The existing multicast extensions for IP/MPLS are point-to-multipoint extensions for RSVP-TE and point-to-multipoint and multipoint-to-multipoint extensions for LDP. These deployment scenarios are discussed in RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*. This feature overview is limited to point-to-multipoint extensions for LDP.

How M-LDP Works

IN THIS SECTION


- [Label Bindings in M-LDP Signaling | 1073](#)
- [M-LDP in PIM-Free MPLS Core | 1074](#)
- [M-LDP in PIM-Enabled MPLS Core | 1076](#)

Label Bindings in M-LDP Signaling

The multipoint extension to LDP uses point-to-multipoint and multipoint-to-multipoint forwarding equivalence class (FEC) elements (defined in RFC 5036, *LDP Specification*) along with capability advertisements, label mapping, and signaling procedures. The FEC elements include the idea of the LSP root, which is an IP address, and an “opaque” value, which is a selector that groups together the leaf nodes sharing the same opaque value. The opaque value is transparent to the intermediate nodes, but has meaning for the LSP root. Every LDP node advertises its local incoming label binding to the upstream LDP node on the shortest path to the root IP address found in the FEC. The upstream node receiving the label bindings creates its own local label and outgoing interfaces. This label allocation process might result in packet replication, if there are multiple outgoing branches. As shown in [Figure 89 on page 1074](#), an LDP node merges

the label bindings for the same opaque value if it finds downstream nodes sharing the same upstream node. This allows for effective building of point-to-multipoint LSPs and label conservation.


Figure 89: Label Bindings in M-LDP Signaling



M-LDP in PIM-Free MPLS Core

Figure 90 on page 1075 shows a scaled-down deployment scenario. Two separate PIM domains are interconnected by a PIM-free core site. The border routers in this core site support PIM on the border interfaces. Further, these border routers collect and distribute the routing information from the adjacent sites to the core network. The edge routers in Site C run BGP for root-node discovery. Interior gateway protocol (IGP) routes cannot be used for ingress discovery because in most cases the forwarding next hop provided by the IGP would not provide information about the ingress device toward the source. M-LDP inband signaling has a one-to-one mapping between the point-to-multipoint LSP and the (S,G) flow. With in-band signaling, PIM messages are directly translated into M-LDP FEC bindings. In contrast, out-of-band signaling is based on manual configuration. One application for M-LDP inband signaling is to carry IPTV multicast traffic in an MPLS backbone.

Figure 90: Sample M-LDP Topology in PIM-Free MPLS Core



Configuration

The configuration statement **mldp-inband-signalling** on the label-edge router (LER) enables PIM to use M-LDP in-band signaling for the upstream neighbors when the LER does not detect a PIM upstream neighbor. Static configuration of the MPLS LSP root is included in the PIM configuration, using policy. This is needed when IBGP is not available in the core site or to override IBGP-based LSP root detection.

For example:

```
protocols {
  pim {
    mldp-inband-signalling {
      policy lsp-mapping-policy-example;
    }
  }
}
```

```
policy-options {
  policy-statement lsp-mapping-policy-example {
    term channel1 {
      from {
        source-address-filter ip-prefix</prefix-length>; #policy filter for channel1
      }
      then {
```



```

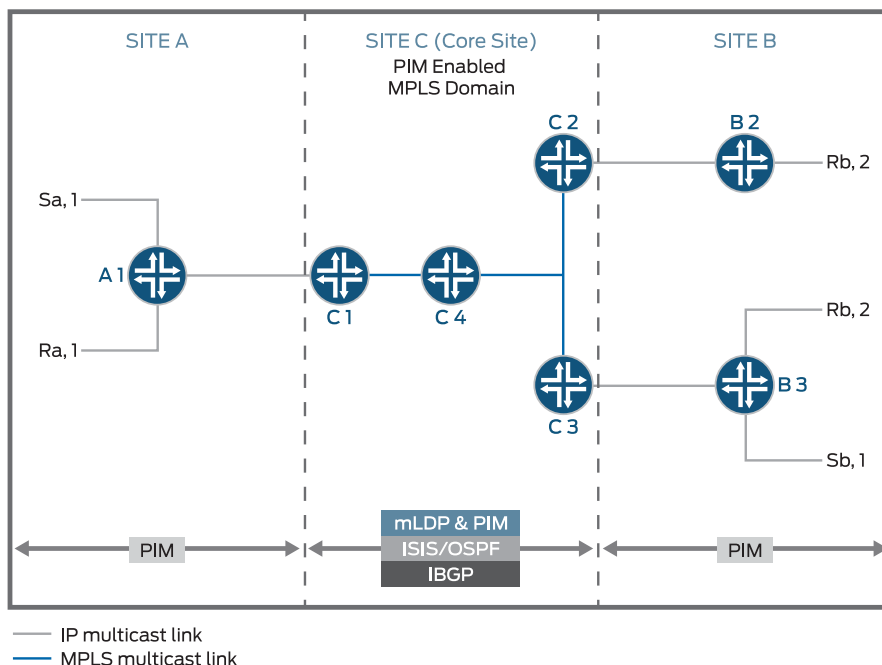
p2mp-lsp-root {
  # Statically configured ingress address of edge
  # used by channel1
  address ip-address;
}
accept;
}
}
}
}
}

```

M-LDP in PIM-Enabled MPLS Core

Starting in Junos OS Release 14.1, in order to migrate existing IPTV services from native IP multicast to MPLS multicast, you need to smoothly transition from PIM to M-LDP point-to-multipoint LSPs with minimal outage. [Figure 91 on page 1076](#) shows a similar M-LDP topology as [Figure 90 on page 1075](#), but with a different scenario. The core is enabled with PIM, with one source streaming all the IPTV channels. The TV channels are sent as ASM streams with each channel identified by its group address. Previously, these channels were streamed on the core as IP streams and signaled using PIM.

Figure 91: Sample M-LDP Topology in PIM-Enabled MPLS Core



By configuring the **mldp-inband-signaling** in this scenario, M-LDP signaling is initiated only when there is no PIM neighbor towards the source. However, because there is always a PIM neighbor towards the source

unless PIM is deactivated on the upstream interfaces of the egress PE, PIM takes precedence over M-LDP and M-LDP does not take effect.

Configuration

To progressively migrate channel by channel to M-LDP MPLS core with few streams using M-LDP upstream and other streams using existing PIM upstream, include the **selected-mldp-egress** configuration statement along with group based filters in the policy filter for M-LDP inband signaling.

NOTE: The M-LDP inband signaling policy filter can include either the **source-address-filter** statement or the **route-filter** statement, or a combination of both.

For example:

```
protocols {
  pim {
    mldp-inband-signalling {
      policy lsp-mapping-policy-example;
    }
  }
}
```

```
policy-options {
  policy-statement lsp-mapping-policy-example {
    term channel1 {
      from {
        source-address-filter ip-prefix</prefix-length>; #policy filter for channel1
      }
      then {
        selected-mldp-egress;
        accept;
      }
    }
    term channel2 {
      from {
        source-address-filter ip-prefix</prefix-length>; #policy filter for channel2
        route-filter ip-prefix</prefix-length>; #policy filter on multicast group address
      }
      then {
        selected-mldp-egress;
        p2mp-lsp-root {
          # Statically configured ingress address of edge
        }
      }
    }
  }
}
```



```

        # used by channel2
        address ip-address;
    }
    accept;
}
}
term channel3 {
    from {
        route-filter ip-prefix</prefix-length>; #policy filter on multicast group address
    }
    then {
        selected-mldp-egress;
        accept;
    }
}
}
}

```

NOTE:

Some of the limitations of the above configuration are as follows:

- The **selected-mldp-egress** statement should be configured only on the LER. Configuring the **selected-mldp-egress** statement on non-egress PIM routers can cause path setup failures.
- When policy changes are made to switch traffic from PIM upstream to M-LDP upstream and vice-versa, packet loss can be expected as break-and-make mechanism is performed at the control plane.

Terminology

The following terms are important for an understanding of M-LDP in-band signaling for multicast traffic.

Point-to-point LSP—An LSP that has one ingress label-switched router (LSR) and one egress LSR.

Multipoint LSP—Either a point-to-multipoint or a multipoint-to-multipoint LSP.

Point-to-multipoint LSP—An LSP that has one ingress LSR and one or more egress LSRs.

Multipoint-to-point LSP—An LSP that has one or more ingress LSRs and one unique egress LSR.

Multipoint-to-multipoint LSP—An LSP that connects a set of nodes, such that traffic sent by any node in the LSP is delivered to all others.

Ingress LSR—An ingress LSR for a particular LSP is an LSR that can send a data packet along the LSP.

Multipoint-to-multipoint LSPs can have multiple ingress LSRs. Point-to-multipoint LSPs have only one, and that node is often referred to as the root node.

Egress LSR—An egress LSR for a particular LSP is an LSR that can remove a data packet from that LSP for further processing. Point-to-point and multipoint-to-point LSPs have only a single egress node.

Point-to-multipoint and multipoint-to-multipoint LSPs can have multiple egress nodes.

Transit LSR—An LSR that has reachability to the root of the multipoint LSP through a directly connected upstream LSR and one or more directly connected downstream LSRs.

Bud LSR—An LSR that is an egress but also has one or more directly connected downstream LSRs.

Leaf node—Either an egress or bud LSR in the context of a point-to-multipoint LSP. In the context of a multipoint-to-multipoint LSP, an LSR is both ingress and egress for the same multipoint-to-multipoint LSP and can also be a bud LSR.

Ingress Join Translation and Pseudo Interface Handling

At the ingress LER, LDP notifies PIM about the (S,G) messages that are received over the in-band signaling. PIM associates each (S,G) message with a pseudo interface. Subsequently, a shortest-path-tree (SPT) join message is initiated toward the source. PIM treats this as a new type of local receiver. When the LSP is torn down, PIM removes this local receiver based on notification from LDP.

Ingress Splicing

LDP provides PIM with a next hop to be associated with each (S,G) entry. PIM installs a PIM (S,G) multicast route with the LDP next hop and other PIM receivers. The next hop is a composite next hop of local receivers + the list of PIM downstream neighbors + a sub-level next hop for the LDP tunnel.

Reverse Path Forwarding

PIM's reverse-path-forwarding (RPF) calculation is performed at the egress node.

PIM performs M-LDP in-band signaling when all of the following conditions are true:

- There are no PIM neighbors toward the source.
- The M-LDP in-band signaling statement is configured.
- The next hop is learned through BGP, or is present in the static mapping (specified in an M-LDP in-band signaling policy).

Otherwise, if LSP root detection fails, PIM retains the (S,G) entry with an RPF state of unresolved.

PIM RPF registers this source address each time unicast routing information changes. Therefore, if the route toward the source changes, the RPF recalculation recurs. BGP protocol next hops toward the source too are monitored for changes in the LSP root. Such changes might cause traffic disruption for short durations.

LSP Root Detection

If the RPF operation detects the need for M-LDP in-band signaling upstream, the LSP root (ingress) is detected. This root is a parameter for LDP LSP signaling.

The root node is detected as follows:

1. If the existing static configuration specifies the source address, the root is taken as given in configuration.
2. A lookup is performed in the unicast routing table. If the source address is found, the protocol next hop toward the source is used as the LSP root.

Prior to Junos OS Release 16.1, M-LDP point-to-multipoint LSP is signaled from an egress to ingress using the root address of the ingress LSR. This root address is reachable through IGP only, thereby confining the M-LDP point-to-multipoint LSP to a single autonomous system. If the root address is not reachable through an IGP, but reachable through BGP, and if that BGP route is recursively resolved over an MPLS LSP, then the point-to-multipoint LSP is not signaled further from that point towards the ingress LSR root address.

There is a need for these non-segmented point-to-multipoint LSPs to be signaled across multiple autonomous systems, which can be used for the following applications:

- Inter-AS MVPN with non-segmented point-to-multipoint LSPs.
- Inter-AS M-LDP inband signaling between client networks connected by an MPLS core network.
- Inter-area MVPN or M-LDP inband signaling with non-segmented point-to-multipoint LSPs (seamless MPLS multicast).

Starting from Junos OS Release 16.1, M-LDP can signal point-to-multipoint LSPs at ASBR or transit or egress when root address is a BGP route which is further recursively resolved over an MPLS LSP.

Egress Join Translation and Pseudo Interface Handling

At the egress LER, PIM notifies LDP of the (S,G) message to be signaled along with the LSP root. PIM creates a pseudo interface as the upstream interface for this (S,G) message. When an (S,G) prune message is received, this association is removed.

Egress Splicing

At the egress node of the core network, where the (S,G) join message from the downstream site is received, this join message is translated to M-LDP in-band signaling parameters and LDP is notified. Further, LSP teardown occurs when the (S,G) entry is lost, when the LSP root changes, or when the (S,G) entry is reachable over a PIM neighbor.

Supported Functionality

For M-LDP in-band signaling, Junos OS supports the following functionality:

- Egress splicing of the PIM next hop with the LDP route
- Ingress splicing of the PIM route with the LDP next hop

- Translation of PIM join messages to LDP point-to-multipoint LSP setup parameters
- Translation of M-LDP in-band LSP parameters to set up PIM join messages
- Statically configured and BGP protocol next hop-based LSP root detection
- PIM (S,G) states in the PIM source-specific multicast (SSM) and anysource multicast (ASM) ranges
- Configuration statements on ingress and egress LERs to enable them to act as edge routers
- IGMP join messages on LERs
- Carrying IPv6 source and group address as opaque information toward an IPv4 root node
- Static configuration to map an IPv6 (S,G) to an IPv4 root address

Unsupported Functionality

For M-LDP in-band signaling, Junos OS does *not* support the following functionality:

- Full support for PIM ASM
- The **mpls lsp point-to-multipoint ping** command with an (S,G) option
- Nonstop active routing (NSR)
- Make-before-break (MBB) for PIM
- IPv6 LSP root addresses (LDP does not support IPv6 LSPs.)
- Neighbor relationship between PIM speakers that are not directly connected
- Graceful restart
- PIM dense mode
- PIM bidirectional mode

LDP Functionality

The PIM (S,G) information is carried as M-LDP opaque type-length-value (TLV) encodings. The point-to-multipoint FEC element consists of the root-node address. In the case of next-generation multicast VPNs (NGEN MVPNs), the point-to-multipoint LSP is identified by the root node address and the LSP ID.

Egress LER Functionality

On the egress LER, PIM triggers LDP with the following information to create a point-to-multipoint LSP:

- Root node
- (S,G)
- Next hop

PIM finds the root node based on the source of the multicast tree. If the root address is configured for this (S,G) entry, the configured address is used as the point-to-multipoint LSP root. Otherwise, the routing table is used to look up the route to the source. If the route to the source of the multicast tree is a

BGP-learned route, PIM retrieves the BGP next hop address and uses it as the root node for the point-to-multipoint LSP.

LDP finds the upstream node based on the root node, allocates a label, and sends the label mapping to the upstream node. LDP does not use penultimate hop popping (PHP) for in-band M-LDP signaling.

If the root addresses for the source of the multicast tree changes, PIM deletes the point-to-multipoint LSP and triggers LDP to create a new point-to-multipoint LSP. When this happens, the outgoing interface list becomes NULL, PIM triggers LDP to delete the point-to-multipoint LSP, and LDP sends a label withdraw message to the upstream node.

Transit LSR Functionality

The transit LSR advertises a label to the upstream LSR toward the source of the point-to-multipoint FEC and installs the necessary forwarding state to forward the packets. The transit LSR can be any M-LDP capable router.

Ingress LER Functionality

On the ingress LER, LDP provides the following information to PIM upon receiving the label mapping:

- (S,G)
- Flood next hop

Then PIM installs the forwarding state. If the new branches are added or deleted, the flood next hop is updated accordingly. If all branches are deleted due to a label being withdrawn, LDP sends updated information to PIM. If there are multiple links between the upstream and downstream neighbors, the point-to-multipoint LSP is not load balanced.

SEE ALSO

| [Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs](#) | 1082

Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs

IN THIS SECTION

- [Requirements](#) | 1083
- [Overview](#) | 1083
- [Configuration](#) | 1084
- [Verification](#) | 1096

This example shows how to configure multipoint LDP (M-LDP) in-band signaling for multicast traffic, as an extension to the Protocol Independent Multicast (PIM) protocol or as a substitute for PIM.

Requirements

This example can be configured using the following hardware and software components:

- Junos OS Release 13.2 or later
- MX Series 5G Universal Routing Platforms or M Series Multiservice Edge Routers for the Provider Edge (PE) Routers
- PTX Series Packet Transport Routers acting as transit label-switched routers
- T Series Core Routers for the Core Routers


NOTE: The PE routers could also be T Series Core Routers but that is not typical. Depending on your scaling requirements, the core routers could also be MX Series 5G Universal Routing Platforms or M Series Multiservice Edge Routers. The Customer Edge (CE) devices could be other routers or switches from Juniper Networks or another vendor.

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

Overview

[“CLI Quick Configuration” on page 1084](#) shows the configuration for all of the devices in [Figure 92 on page 1084](#). The section [“Step-by-Step Procedure” on page 1089](#) describes the steps on Device EgressPE.

Figure 92: M-LDP In-Band Signaling for Point-to-Multipoint LSPs Example Topology



Configuration

CLI Quick Configuration

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

Device src1

```
set logical-systems src1 interfaces fe-1/2/0 unit 0 family inet address 1.2.7.7/24
set logical-systems src1 interfaces lo0 unit 0 family inet address 1.1.1.7/32
set logical-systems src1 protocols ospf area 0.0.0.0 interface all
```

Device IngressPE

```
set interfaces so-0/1/2 unit 0 family inet address 192.168.93.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.2.3.2/24
set interfaces fe-1/2/0 unit 0 family mpls
```



```

set interfaces fe-1/2/1 unit 0 family inet address 1.2.5.2/24
set interfaces fe-1/2/2 unit 0 family inet address 1.2.6.2/24
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/2/3 unit 0 family inet address 1.2.7.2/24
set interfaces fe-1/3/1 unit 0 family inet address 192.168.219.9/28
set interfaces lo0 unit 0 family inet address 1.1.1.2/32
set protocols igmp interface fe-1/2/1.0 version 3
set protocols igmp interface fe-1/2/1.0 static group 232.1.1.1 source 192.168.219.11
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.2
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp neighbor 1.1.1.3
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols bgp group ibgp neighbor 1.1.1.1
set protocols ospf area 0.0.0.0 interface all
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim rp static address 1.1.1.5
set protocols pim interface fe-1/3/1.0
set protocols pim interface lo0.0
set protocols pim interface fe-1/2/0.21
set protocols pim interface fe-1/2/3.0
set protocols pim interface fe-1/2/1.0
set protocols pim interface so-0/1/2.0
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
    orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
    orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 1.1.1.7/32 orlonger
set policy-options policy-statement mldppim-ex term A from source-address-filter 1.2.7.0/24 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set routing-options autonomous-system 64510

```

Device EgressPE


```

set interfaces so-0/1/3 unit 0 point-to-point
set interfaces so-0/1/3 unit 0 family inet address 192.168.92.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.1.3.1/24
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 1.1.4.1/24
set interfaces fe-1/2/2 unit 0 family inet address 1.1.6.1/24
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/3/0 unit 0 family inet address 192.168.209.9/28
set interfaces lo0 unit 0 family inet address 1.1.1.1/32
set routing-options autonomous-system 64510
set protocols igmp interface fe-1/3/0.0 version 3
set protocols igmp interface fe-1/3/0.0 static group 232.1.1.1 group-count 3
set protocols igmp interface fe-1/3/0.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface fe-1/3/0.0 static group 227.1.1.1
set protocols igmp interface so-0/1/3.0 version 3
set protocols igmp interface so-0/1/3.0 static group 232.1.1.1 group-count 2
set protocols igmp interface so-0/1/3.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface so-0/1/3.0 static group 232.2.2.2 source 1.2.7.7
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.1
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols msdp local-address 1.1.1.1
set protocols msdp peer 1.1.1.5
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim rp local address 1.1.1.1
set protocols pim rp local group-ranges 227.0.0.0/8
set protocols pim rp static address 1.1.1.4
set protocols pim rp static address 1.2.7.7 group-ranges 226.0.0.0/8
set protocols pim interface lo0.0
set protocols pim interface fe-1/3/0.0
set protocols pim interface fe-1/2/0.0
set protocols pim interface fe-1/2/1.0
set protocols pim interface so-0/1/3.0

```



```

set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
  orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter 1.2.7.0/24 orlonger
set policy-options policy-statement mldppim-ex term A then accept

```

Device p6

```

set interfaces fe-1/2/0 unit 0 family inet address 1.1.6.6/24
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 1.2.6.6/24
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.6/32
set interfaces lo0 unit 0 family mpls
set protocols ospf area 0.0.0.0 interface all
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/1.0
set protocols ldp interface lo0.0
set protocols ldp p2mp

```

Device pr3

```

set interfaces ge-0/3/1 unit 0 family inet address 192.168.215.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.1.3.3/24
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 1.2.3.3/24
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.3/32
set protocols igmp interface ge-0/3/1.0 version 3
set protocols igmp interface ge-0/3/1.0 static group 232.1.1.2 source 192.168.219.11
set protocols igmp interface ge-0/3/1.0 static group 232.2.2.2 source 1.2.7.7
set protocols bgp group ibgp local-address 1.1.1.3
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 1.1.1.2

```



```

set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fe-1/2/1.0 metric 2
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/1.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface fe-0/3/1.0
set protocols pim interface lo0.0
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
  orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.0.0/24
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 192.168.219.11/32
  orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter 1.2.7.7/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 64510

```

Device pr4

```

set interfaces ge-0/3/0 unit 0 family inet address 192.168.207.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.1.4.4/24
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 1.1.1.4/32
set protocols igmp interface ge-0/3/0.0 version 3
set protocols igmp interface ge-0/3/0.0 static group 232.1.1.2 source 192.168.219.11
set protocols igmp interface ge-0/3/0.0 static group 225.1.1.1
set protocols bgp group ibgp local-address 1.1.1.4
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols msdp local-address 1.1.1.4
set protocols msdp peer 1.1.1.5
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 1.1.1.4

```



```

set protocols pim interface ge-0/3/0.0
set protocols pim interface lo0.0
set protocols pim interface fe-1/2/0.0
set routing-options autonomous-system 64510

```

Device pr5

```

set interfaces fe-1/2/0 unit 0 family inet address 1.2.5.5/24
set interfaces lo0 unit 0 family inet address 1.1.1.5/24
set protocols igmp interface lo0.0 version 3
set protocols igmp interface lo0.0 static group 232.1.1.1 source 192.168.219.11
set protocols msdp local-address 1.1.1.5
set protocols msdp peer 1.1.1.4
set protocols msdp peer 1.1.1.1
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 1.1.1.5
set protocols pim interface all

```

Step-by-Step Procedure

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

To configure Device EgressPE:

1. Configure the interfaces.

Enable MPLS on the core-facing interfaces. On the egress next hops, you do not need to enable MPLS.

```

[edit interfaces]
user@EgressPE# set fe-1/2/0 unit 0 family inet address 1.1.3.1/24
user@EgressPE# set fe-1/2/0 unit 0 family mpls
user@EgressPE# set fe-1/2/2 unit 0 family inet address 1.1.6.1/24
user@EgressPE# set fe-1/2/2 unit 0 family mpls
user@EgressPE# set so-0/1/3 unit 0 point-to-point
user@EgressPE# set so-0/1/3 unit 0 family inet address 192.168.92.9/28
user@EgressPE# set fe-1/2/1 unit 0 family inet address 1.1.4.1/24
user@EgressPE# set fe-1/3/0 unit 0 family inet address 192.168.209.9/28
user@EgressPE# set lo0 unit 0 family inet address 1.1.1.1/32

```


2. Configure IGMP on the egress interfaces.

For testing purposes, this example includes static group and source addresses.

```
[edit protocols igmp]
user@EgressPE# set interface fe-1/3/0.0 version 3
user@EgressPE# set interface fe-1/3/0.0 static group 232.1.1.1 group-count 3
user@EgressPE# set interface fe-1/3/0.0 static group 232.1.1.1 source 192.168.219.11
user@EgressPE# set interface fe-1/3/0.0 static group 227.1.1.1
user@EgressPE# set interface so-0/1/3.0 version 3
user@EgressPE# set interface so-0/1/3.0 static group 232.1.1.1 group-count 2
user@EgressPE# set interface so-0/1/3.0 static group 232.1.1.1 source 192.168.219.11
user@EgressPE# set interface so-0/1/3.0 static group 232.2.2.2 source 1.2.7.7
```

3. Configure MPLS on the core-facing interfaces.

```
[edit protocols mpls]
user@EgressPE# set interface fe-1/2/0.0
user@EgressPE# set interface fe-1/2/2.0
```

4. Configure BGP.

BGP is a policy-driven protocol, so also configure and apply any needed routing policies.

For example, you might want to export static routes into BGP.

```
[edit protocols bgp group ibgp]
user@EgressPE# set type internal
user@EgressPE# set local-address 1.1.1.1
user@EgressPE# set family inet any
user@EgressPE# set neighbor 1.1.1.2
```

5. (Optional) Configure an MSDP peer connection with Device pr5 in order to interconnect the disparate PIM domains, thus enabling redundant RPs.

```
[edit protocols msdp]
user@EgressPE# set local-address 1.1.1.1
user@EgressPE# set peer 1.1.1.5
```

6. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
```



```

user@EgressPE# set interface all
user@EgressPE# set interface fxp0.0 disable

```

7. Configure LDP on the core-facing interfaces and on the loopback interface.

```

[edit protocols ldp]
user@EgressPE# set interface fe-1/2/0.0
user@EgressPE# set interface fe-1/2/2.0
user@EgressPE# set interface lo0.0

```

8. Enable point-to-multipoint MPLS LSPs.

```

[edit protocols ldp]
user@EgressPE# set p2mp

```

9. Configure PIM on the downstream interfaces.

```

[edit protocols pim]
user@EgressPE# set interface lo0.0
user@EgressPE# set interface fe-1/3/0.0
user@EgressPE# set interface fe-1/2/1.0
user@EgressPE# set interface so-0/1/3.0

```

10. Configure the RP settings because this device serves as the PIM rendezvous point (RP).

```

[edit protocols pim]
user@EgressPE# set rp local address 1.1.1.1
user@EgressPE# set rp local group-ranges 227.0.0.0/8
user@EgressPE# set rp static address 1.1.1.4
user@EgressPE# set rp static address 1.2.7.7 group-ranges 226.0.0.0/8

```

11. Enable M-LDP in-band signaling and set the associated policy.

```

[edit protocols pim]
user@EgressPE# set mldp-inband-signalling policy mldppim-ex

```

12. Configure the routing policy that specifies the root address for the point-to-multipoint LSP and the associated source addresses.


```
[edit policy-options policy-statement mldppim-ex]
user@EgressPE# set term B from source-address-filter 192.168.0.0/24 orlonger
user@EgressPE# set term B from source-address-filter 192.168.219.11/32 orlonger
user@EgressPE# set term B then p2mp-lsp-root address 1.1.1.2
user@EgressPE# set term B then accept
user@EgressPE# set term A from source-address-filter 1.2.7.0/24 orlonger
user@EgressPE# set term A then accept
```

13. Configure the autonomous system (AS) ID.

```
[edit routing-options]
user@EgressPE# set autonomous-system 64510
```

Results

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

Device EgressPE

```
user@EgressPE# show interfaces
so-0/1/3 {
  unit 0 {
    point-to-point;
    family inet {
      address 192.168.92.9/28;
    }
  }
}
fe-1/2/0 {
  unit 0 {
    family inet {
      address 1.1.3.1/24;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    family inet {
```



```

        address 1.1.4.1/24;
    }
}
fe-1/2/2 {
    unit 0 {
        family inet {
            address 1.1.6.1/24;
        }
        family mpls;
    }
}
fe-1/3/0 {
    unit 0 {
        family inet {
            address 192.168.209.9/28;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 1.1.1.1/32;
        }
    }
}

```

user@EgressPE# **show protocols**

```

igmp {
    interface fe-1/3/0.0 {
        version 3;
        static {
            group 232.1.1.1 {
                group-count 3;
                source 192.168.219.11;
            }
            group 227.1.1.1;
        }
    }
}
interface so-0/1/3.0 {
    version 3;
}

```



```

static {
    group 232.1.1.1 {
        group-count 2;
        source 192.168.219.11;
    }
    group 232.2.2.2 {
        source 1.2.7.7;
    }
}
}
}
mpls {
    interface fe-1/2/0.0;
    interface fe-1/2/2.0;
}
bgp {
    group ibgp {
        type internal;
        local-address 1.1.1.1;
        family inet {
            any;
        }
        neighbor 1.1.1.2;
    }
}
msdp {
    local-address 1.1.1.1;
    peer 1.1.1.5;
}
ospf {
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface fe-1/2/0.0;
    interface fe-1/2/2.0;
    interface lo0.0;
    p2mp;
}
pim {

```



```

mldp-inband-signalling {
    policy mldppim-ex;
}
rp {
    local {
        address 1.1.1.1;
        group-ranges {
            227.0.0.0/8;
        }
    }
    static {
        address 1.1.1.4;
        address 1.2.7.7 {
            group-ranges {
                226.0.0.0/8;
            }
        }
    }
}
interface lo0.0;
interface fe-1/3/0.0;
interface fe-1/2/0.0;
interface fe-1/2/1.0;
interface so-0/1/3.0;
}

```

```

user@EgressPE# show policy-options
policy-statement mldppim-ex {
    term B {
        from {
            source-address-filter 192.168.0.0/24 orlonger;
            source-address-filter 192.168.219.11/32 orlonger;
        }
        then {
            p2mp-lsp-root {
                address 1.1.1.2;
            }
            accept;
        }
    }
    term A {
        from {
            source-address-filter 1.2.7.0/24 orlonger;
        }
    }
}

```



```

    then accept;
  }
}

```

```

user@EgressPE# show routing-options
autonomous-system 64510;

```

Similarly, configure the other egress devices.

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

Verification

IN THIS SECTION

- [Checking the PIM Join States | 1096](#)
- [Checking the PIM Sources | 1101](#)
- [Checking the LDP Database | 1103](#)
- [Looking Up the Route Information for the MPLS Label | 1107](#)
- [Checking the LDP Traffic Statistics | 1108](#)

Confirm that the configuration is working properly.

Checking the PIM Join States

Purpose

Display information about PIM join states to verify the M-LDP in-band upstream and downstream details. On the ingress device, the **show pim join extensive** command displays **Pseudo-MLDP** for the downstream interface. On the egress, the **show pim join extensive** command displays **Pseudo-MLDP** for the upstream interface.

Action

From operational mode, enter the **show pim join extensive** command.

```
user@IngressPE> show pim join extensive
```

```

Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```



```

Group: 232.1.1.1
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 1d 23:00:12
  Downstream neighbors:
    Interface: Pseudo-MLDP
    Interface: fe-1/2/1.0
      1.2.5.2 State: Join Flags: S   Timeout: Infinity
      Uptime: 1d 23:00:12 Time since last Join: 1d 23:00:12

Group: 232.1.1.2
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 1d 22:59:59
  Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.1.1.3
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 1d 22:07:31
  Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.2.2.2
  Source: 1.2.7.7
  Flags: sparse,spt
  Upstream interface: fe-1/2/3.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 1d 22:59:59

```



```
Downstream neighbors:
  Interface: Pseudo-MLDP
```

```
user@EgressPE> show pim join extensive
```

```
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
  Source: *
  RP: 1.1.1.1
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 1d 23:14:21
  Downstream neighbors:
    Interface: fe-1/3/0.0
      192.168.209.9 State: Join Flags: SRW Timeout: Infinity
      Uptime: 1d 23:14:21 Time since last Join: 1d 20:12:35

Group: 232.1.1.1
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 1d 23:14:22
  Downstream neighbors:
    Interface: so-0/1/3.0
      192.168.92.9 State: Join Flags: S Timeout: Infinity
      Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35
    Downstream neighbors:
      Interface: fe-1/3/0.0
        192.168.209.9 State: Join Flags: S Timeout: Infinity
        Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35

Group: 232.1.1.2
  Source: 192.168.219.11
  Flags: sparse,spt
```



```

Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 23:14:22
Downstream neighbors:
  Interface: so-0/1/3.0
    192.168.92.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35
Downstream neighbors:
  Interface: fe-1/2/1.0
    1.1.4.4 State: Join Flags: S Timeout: 198
    Uptime: 1d 22:59:59 Time since last Join: 00:00:12
Downstream neighbors:
  Interface: fe-1/3/0.0
    192.168.209.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35

```

```

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 20:12:35
Downstream neighbors:
  Interface: fe-1/3/0.0
    192.168.209.9 State: Join Flags: S   Timeout: Infinity
    Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35

```

```

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 20:12:35
Downstream neighbors:
  Interface: so-0/1/3.0

```



```
192.168.92.9 State: Join Flags: S   Timeout: Infinity
Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35
```

user@pr3> **show pim join extensive**

```
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 232.1.1.2
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 1d 20:14:40
  Downstream neighbors:
    Interface: Pseudo-GMP
    ge-0/3/1.0
```

```
Group: 232.2.2.2
  Source: 1.2.7.7
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 1d 20:14:40
  Downstream neighbors:
    Interface: Pseudo-GMP
    ge-0/3/1.0
```

user@pr4> **show pim join extensive**

```
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
  Source: *
```



```

RP: 1.1.1.4
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 1d 23:13:43
Downstream neighbors:
  Interface: ge-0/3/0.0
    192.168.207.9 State: Join Flags: SRW Timeout: Infinity
    Uptime: 1d 23:13:43 Time since last Join: 1d 23:13:43

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/2/0.0
Upstream neighbor: 1.1.4.1
Upstream state: Local RP, Join to Source
Keepalive timeout: 0
Uptime: 1d 23:13:43
Downstream neighbors:
  Interface: ge-0/3/0.0
    192.168.207.9 State: Join Flags: S Timeout: Infinity
    Uptime: 1d 23:13:43 Time since last Join: 1d 23:13:43

```

user@pr5> **show pim join extensive**

```

ge-0/3/1.0

Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

Checking the PIM Sources

Purpose

Verify that the PIM sources have the expected M-LDP in-band upstream and downstream details.

Action

From operational mode, enter the **show pim source** command.

user@IngressPE> **show pim source**

```
Instance: PIM.master Family: INET

Source 1.1.1.1
  Prefix 1.1.1.1/32
  Upstream interface Local
  Upstream neighbor Local

Source 1.2.7.7
  Prefix 1.2.7.0/24
  Upstream protocol MLDP
  Upstream interface Pseudo MLDP
  Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11
  Prefix 192.168.219.0/28
  Upstream protocol MLDP
  Upstream interface Pseudo MLDP
  Upstream neighbor MLDP LSP root <1.1.1.2>
```

user@EgressPE> **show pim source**

```
Instance: PIM.master Family: INET

Source 1.2.7.7
  Prefix 1.2.7.0/24
  Upstream interface fe-1/2/3.0
  Upstream neighbor 1.2.7.2

Source 1.2.7.7
  Prefix 1.2.7.0/24
  Upstream interface fe-1/2/3.0
  Upstream neighbor Direct

Source 192.168.219.11
  Prefix 192.168.219.0/28
  Upstream interface fe-1/3/1.0
  Upstream neighbor 192.168.219.9
```



```
Source 192.168.219.11
  Prefix 192.168.219.0/28
  Upstream interface fe-1/3/1.0
  Upstream neighbor Direct
```

user@pr3> **show pim source**

```
Instance: PIM.master Family: INET

Source 1.2.7.7
  Prefix 1.2.7.0/24
  Upstream protocol MLDP
  Upstream interface Pseudo MLDP
  Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11
  Prefix 192.168.219.0/28
  Upstream protocol MLDP
  Upstream interface Pseudo MLDP
  Upstream neighbor MLDP LSP root <1.1.1.2>
```

user@pr4> **show pim source**

```
Instance: PIM.master Family: INET

Source 1.1.1.4
  Prefix 1.1.1.4/32
  Upstream interface Local
  Upstream neighbor Local

Source 192.168.219.11
  Prefix 192.168.219.0/28
  Upstream interface fe-1/2/0.0
  Upstream neighbor 1.1.4.1
```

Checking the LDP Database

Purpose

Make sure that the `show ldp database` command displays the expected root-to-(S,G) bindings.

Action

user@IngressPE> `show ldp database`

Input label database, 10.255.2.227:0--1.1.1.3:0

Label	Prefix
300096	1.1.1.2/32
3	1.1.1.3/32
299856	1.1.1.6/32
299776	10.255.2.227/32

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

Label	Prefix
300144	1.1.1.2/32
299776	1.1.1.3/32
299856	1.1.1.6/32
3	10.255.2.227/32

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

Label	Prefix
299936	1.1.1.2/32
299792	1.1.1.3/32
3	1.1.1.6/32
299776	10.255.2.227/32

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

Label	Prefix
300144	1.1.1.2/32
299776	1.1.1.3/32
299856	1.1.1.6/32
3	10.255.2.227/32
300432	P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
300288	P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
300160	P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
300480	P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11

user@EgressPE> `show ldp database`

Input label database, 1.1.1.2:0--1.1.1.3:0

Label	Prefix
-------	--------


```

300096      1.1.1.2/32
      3      1.1.1.3/32
299856      1.1.1.6/32
299776      10.255.2.227/32
300144      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
300128      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

```

Output label database, 1.1.1.2:0--1.1.1.3:0

```

Label      Prefix
      3      1.1.1.2/32
299776      1.1.1.3/32
299808      1.1.1.6/32
299792      10.255.2.227/32

```

Input label database, 1.1.1.2:0--1.1.1.6:0

```

Label      Prefix
299936      1.1.1.2/32
299792      1.1.1.3/32
      3      1.1.1.6/32
299776      10.255.2.227/32
300128      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299984      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299952      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
300176      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
300192      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

```

Output label database, 1.1.1.2:0--1.1.1.6:0

```

Label      Prefix
      3      1.1.1.2/32
299776      1.1.1.3/32
299808      1.1.1.6/32
299792      10.255.2.227/32

```

logical-system: default

Input label database, 10.255.2.227:0--1.1.1.3:0

```

Label      Prefix
300096      1.1.1.2/32
      3      1.1.1.3/32
299856      1.1.1.6/32
299776      10.255.2.227/32

```

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

Label	Prefix
300144	1.1.1.2/32
299776	1.1.1.3/32
299856	1.1.1.6/32
3	10.255.2.227/32

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

Label	Prefix
299936	1.1.1.2/32
299792	1.1.1.3/32
3	1.1.1.6/32
299776	10.255.2.227/32

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

Label	Prefix
300144	1.1.1.2/32
299776	1.1.1.3/32
299856	1.1.1.6/32
3	10.255.2.227/32
300432	P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
300288	P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
300160	P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
300480	P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
300496	P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

user@p6> show ldp database

Input label database, 1.1.1.6:0--1.1.1.2:0

Label	Prefix
3	1.1.1.2/32
299776	1.1.1.3/32
299808	1.1.1.6/32

Output label database, 1.1.1.6:0--1.1.1.2:0

Label	Prefix
299776	1.1.1.2/32
299792	1.1.1.3/32
3	1.1.1.6/32

user@pr3> show ldp database


```
Input label database, 1.1.1.3:0--1.1.1.2:0
```

Label	Prefix
3	1.1.1.2/32
299776	1.1.1.3/32
299808	1.1.1.6/32
299792	10.255.2.227/32

```
Output label database, 1.1.1.3:0--1.1.1.2:0
```

Label	Prefix
300096	1.1.1.2/32
3	1.1.1.3/32
299856	1.1.1.6/32
299776	10.255.2.227/32
300144	P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
300128	P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

```
Input label database, 1.1.1.3:0--10.255.2.227:0
```

Label	Prefix
300144	1.1.1.2/32
299776	1.1.1.3/32
299856	1.1.1.6/32
3	10.255.2.227/32

```
Output label database, 1.1.1.3:0--10.255.2.227:0
```

Label	Prefix
300096	1.1.1.2/32
3	1.1.1.3/32
299856	1.1.1.6/32
299776	10.255.2.227/32

Looking Up the Route Information for the MPLS Label

Purpose

Display the point-to-multipoint FEC information.

Action

```
user@EgressPE> show route label 299808 detail
```

```
mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
299808 (1 entry, 1 announced)
```



```

*LDP      Preference: 9
          Next hop type: Flood
          Address: 0x931922c
          Next-hop reference count: 3
          Next hop type: Router, Next hop index: 1109
          Address: 0x9318b0c
          Next-hop reference count: 2
          Next hop: via so-0/1/3.0
          Label operation: Pop
          Next hop type: Router, Next hop index: 1110
          Address: 0x93191e0
          Next-hop reference count: 2
          Next hop: 192.168.209.11 via fe-1/3/0.0
          Label operation: Pop
          State: **Active Int AckRequest>
          Local AS:      10
          Age: 13:08:15   Metric: 1
          Validation State: unverified
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
          FECs bound to route: P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src:
192.168.219.11

```

Checking the LDP Traffic Statistics

Purpose

Monitor the data traffic statistics for the point-to-multipoint LSP.

Action

```
user@EgressPE> show ldp traffic-statistics p2mp
```

P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)	Nexthop	Packets	Bytes
Shared			
1.1.1.2:232.2.2.2,1.2.7.7	so-0/1/3.0	0	0
No			
1.1.1.2:232.1.1.1,192.168.219.11	so-0/1/3.0	0	0
No			
	fe-1/3/0.0	0	0
No			

1.1.1.2:232.1.1.2,192.168.219.11	so-0/1/3.0	0	0
No			
	fe-1/3/0.0	0	0
No			
	lt-1/2/0.14	0	0
No			
1.1.1.2:232.1.1.3,192.168.219.11	fe-1/3/0.0	0	0
No			
1.1.1.2:ff3e::1:2,abcd::1:2:7:7	fe-1/3/0.0	0	0
No			

SEE ALSO

[Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs | 1072](#)

Release History Table

Release	Description
16.1	Starting from Junos OS Release 16.1, M-LDP can signal point-to-multipoint LSPs at ASBR or transit or egress when root address is a BGP route which is further recursively resolved over an MPLS LSP.
14.1	Starting in Junos OS Release 14.1, in order to migrate existing IPTV services from native IP multicast to MPLS multicast, you need to smoothly transition from PIM to M-LDP point-to-multipoint LSPs with minimal outage.

RELATED DOCUMENTATION

RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems

Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs

LDP Mapping Server for Interoperability of Segment Routing with LDP Overview

IN THIS SECTION

- [Interoperability of Segment Routing with LDP Using OSPF | 1110](#)
- [Interoperability of Segment Routing with LDP Using ISIS | 1112](#)

In an LDP network with gradual deployment of segment routing, there can be islands of devices that support either only LDP, or only segment routing. For the devices to interwork, the LDP mapping server feature is required to be configured on any device in the segment routing network.


The LDP mapping server feature is implemented using either OSPF or ISIS.

Interoperability of Segment Routing with LDP Using OSPF

To implement interoperability of segment routing with LDP using OSPF, an extended prefix link-state advertisement (LSA) with Range type, length, and value (TLV) for all the LDP prefixes is generated, and mapping routes corresponding to the prefix is installed in the inet.3 and mpls.0 routing tables.

[Figure 93 on page 1110](#) is a simple LDP network topology used to illustrate the interoperability of segment routing devices with LDP devices using OSPF. The topology has six devices (Devices R1, R2, R3, R4, R5, and R6) with LDP-to-segment routing migration.

Figure 93: Sample LDP Topology with Interoperability of Segment Routing with LDP Using OSPF



In the above topology, Devices R1, R2, and R3 are capable of only segment routing, Devices R5 and R6 are capable of only LDP, and Device R4 supports both LDP and segment routing. Here, Device R1 cannot interwork with Device R6 because of interoperability issues.

To enable interoperability between the LDP-capable devices and segment routing devices, any one interface of the device in the segment routing network segment is configured as the LDP mapping server. Currently, the mapping server configures prefixes under the `[edit routing-options source-packet-routing]` hierarchy level. With this feature, the LDP mapping server configuration is applied under the `[edit protocols ospf]`

hierarchy level, where a new extended prefix LSA with range TLV for all LDP prefixes is advertised by OSPF. The device capable of segment routing analyze the extended prefix range TLV and mapping routes corresponding to the prefix are installed in the inet.3 and mpls.0 routing tables.

For example, in [Figure 93 on page 1110](#), if Device R2 (in the segment routing network) is the LDP mapping server, the following configuration is included:

```
[edit routing-options]
user@R2# set source-packet-routing mapping-server-entry mapping-server-name prefix-segment-range prefix-range
start-prefix loopback-address
user@R2# set source-packet-routing mapping-server-entry mp1 prefix-segment-range rg1 start-index 5
user@R2# set source-packet-routing mapping-server-entry mp1 prefix-segment-range rg1 size 1
```

NOTE: The IP address used as the **start-prefix** is the loopback address of the device in the LDP network (Device R5, in this example).

```
[edit protocols]
user@R2# set ospf source-packet-routing mapping-server mapping-server-name
```

When the LDP mapping server configuration is committed on Device R2, the extended prefix range TLV is flooded across the OSPF area. The devices capable of segment routing (Devices R1, R2, and R3) install OSPF segment routing routes for the specified loopback address with a segment ID (SID) index. The SID index is also updated in the mpls.0 routing table by the segment routing devices.

Starting in Junos OS Release 19.1R1, segment routing-LDP border router can stitch segment routing traffic to LDP next-hop and vice versa.

Unsupported Features and Functionality for Interoperability of Segment Routing with LDP using OSPF


- Prefix conflict is only detected at the source configuration. When there is a prefix range conflict, the prefix SID from the lower router ID prevails. In such cases, a system log error message—**RPD_OSPF_PFX_SID_RANGE_CONFLICT**—is generated.
- IPv6 prefixes are not supported.
- Flooding of the OSPF Extended Prefix Opaque LSA generated by the segment routing mapping server for autonomous systems (ASs) is not supported.
- Inter-area LDP mapping server functionality is not supported.
- ABR functionality of Extended Prefix Opaque LSA is not supported.
- ASBR functionality of Extended Prefix Opaque LSA is not supported.
- Segment routing mapping server Preference TLV is not supported.

Interoperability of Segment Routing with LDP Using ISIS

To implement interoperability of segment routing with LDP using ISIS, a server-client configuration is required under protocols ISIS and LDP, respectively, and routes from the inet.3 or inet.0 routing tables are used for stitching of segment routing LSP with an LDP LSP and vice-versa.

Figure 94 on page 1112 is a simple LDP network topology used to illustrate the interoperability of segment routing devices with LDP devices using an LDP mapping server-client feature. The topology has four provider edge (PE) devices (Devices PE1, PE2, PE3, and PE4) and four provider (P) devices (Devices P5, P6, P7, and P8).

Figure 94: Sample LDP Topology with Interoperability of Segment Routing with LDP Using ISIS



Devices PE3, PE4, P6, P7 and P8 are the LDP capable devices. Devices PE1, PE2, P5 and P6 are capable of segment routing with segment routing global block (SRGB) value of 100 and 200, and node segment IDs (SIDs) value of 101, 102, 105 and 106, respectively.

For a service flow to be tunneled to-and-from Device PE3 and Device PE1 using a continuous MPLS tunnel, the islands of devices supporting segment routing and LDP must interoperate.

LDP Mapping Client Functionality (LDP to Segment Routing)

The LDP client functionality is the LDP-to-segment routing mapping, that is the right-to-left traffic flow in Figure 94 on page 1112. On Device P6, an LDP egress policy is configured to advertise all node SIDs and prefix SIDs from the segment routing network on the left. As a result, on Device P6, LDP advertises Devices PE1, PE2 and P5 as the egress FEC label bindings to Device P7.

Device PE3 has learned a service route with Device PE1 as the protocol next hop. Device PE3 has an LDP label binding from the P8 next hop for the PE1 FEC. As a result, Device PE3 sends its service packet to Device P8 as per classic LDP behavior. Device P8 has an LDP label binding from its P7 next hop for the PE1 FEC, therefore Device P8 forwards to Device P7 as per classic LDP behavior.

Device P7 has an LDP label binding from its P6 next hop the PE1 FEC, as a result, Device P7 forwards to Device P6 as per classic LDP behavior.

Device P6 that is acting as an LDP egress for the PE1 FEC, stitches and swaps the incoming egress LDP label for the PE1 FEC with an equivalent segment routing node SID (101 in this example) to forward the traffic to Device P5.

Device P5 pops 101 SID assuming that Device PE1 advertised its node segment 101 with the penultimate-pop flag set, and then forwards traffic to Device PE1. Device PE1 receives the tunneled packet and processes the service label.

As a result, the end-to-end MPLS tunnel is built from an LDP LSP from Device PE3 to Device P6, and the related node segment from Device P6 to Device PE1.

LDP Mapping Server Functionality (Segment Routing to LDP)

The LDP server functionality is the mapping of segment routing to LDP, that is the left-to-right traffic flow in [Figure 94 on page 1112](#). On Device P6 the mapping server prefixes configuration is included under the **[edit routing-options source-packet-routing]** hierarchy level. When the configuration is applied under the specific IGP, the label binding type, length, and value (TLV) for all the LDP FEC-label bindings from the LDP network are advertised as inet.3 LDP routes.

Here, Device P6 acts as a Segment Routing Mapping Server (SRMS) and advertises the following mappings – (P7, 107), (P8, 108), (PE3, 103), (PE4, 104), and (P7, 107). If segment routing was supported on Device PE3, the node SID 103 would have been configured on Device PE3. Because Device PE3 does not support segment routing, the policy is configured at the SRMS on Device P6, and Device P6 is responsible for advertising the mappings.

These mapping server advertisements are only understood by the segment routing devices. The segment routing devices install the related node SIDs in the MPLS data plane exactly how the node segments had been advertised by the nodes themselves. For instance, Device PE1 installs the node SID 103 with P5 next hop exactly as if Device PE3 had advertised SID 103.

Device PE1 has a service route with PE3 as its protocol next hop. Device PE1 has a node segment for that IGP route – 103 with P5 next hop. As a result, Device PE1 sends its service packet to Device P5 with two labels – the bottom label, which is the service label, and the top label, which is SID 103. Device P5 swaps 103 for 103 and forwards to Device P6. The next-hop for Device P6 is the IGP route PE3, which is not capable of segment routing. (Device P7 does not advertise the segment routing capability). However, Device P6 has an LDP label binding from that next hop for the same FEC (for example, LDP label 1037). As a result, on Device P6, the IGP swaps 103 for 1037 and forwards to Device P7.

Device P7 swaps this label with the LDP-label received from Device P8, and then forwards it to Device P8. The LDP label is popped by Device P8 and forwarded to Device PE3.

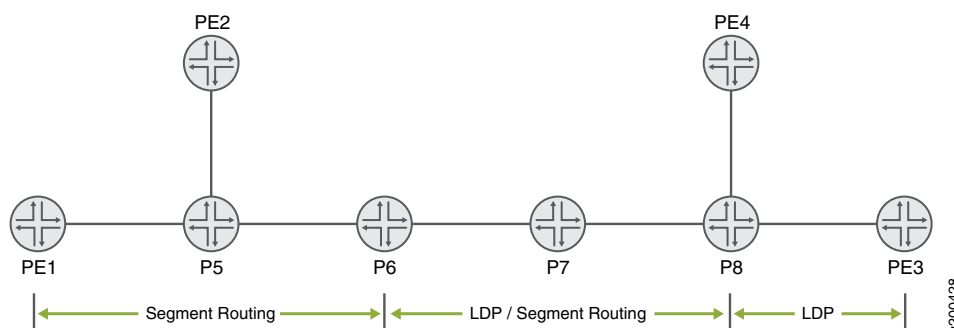
Device PE3 receives the tunneled packet and processes the service label. The end-to-end MPLS tunnel is built from a segment routing node from Devices PE1 to P6, and an LDP LSP from Devices P6 to PE3.

Segment Routing to LDP Stitching

When the IGP segment routing LSP's IP next hop does not support segment routing, the IGP looks at the inet.3 routing table to see if there is an LDP LSP to the same prefix. If the LDP LSP is present, the IGP stitches the segment routing LSP to the LDP LSP by programming a MPLS transit route that swaps the segment routing label with the LDP label to switch traffic from segment routing domain to the LDP domain.

Figure 95 on page 1114 illustrates the stitching of segment routing and LDP LSPs for enabling interoperability.

Figure 95: Stitching Segment Routing and LDP LSPs



In the topology, Device PE3 is LDP-capable and does not support segment routing. The mapping server in the segment routing domain can advertise label binding TLV for devices P7, P8 and PE4. In such a scenario, Device PE1 can have both prefix SID and remote label binding TLV and SID to reach Device PE4. However, Device PE1 prefers prefix SID over remote label binding TLV while programming its ingress segment routing route for Device PE4. As a result, Device PE1 uses the segment routing LSP end-to-end to send traffic to Device PE4, and uses the segment routing-to-LDP stitching while sending traffic to Device PE3.

Unsupported Features and Functionality for Interoperability of Segment Routing with LDP using ISIS

- Penultimate-hop popping behaviour for label binding TLV is not supported.
- Advertising of range of prefixes in label binding TLV is not supported.
- Segment Routing Conflict Resolution is not supported.
- LDP traffic statistics does not work.
- Nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) is not supported.
- ISIS inter-level is not supported.
- RFC 7794, *IS-IS Prefix Attributes for Extended IPv4* is not supported.
- Redistributing LDP route as a prefix-sid at the stitching node is not supported.

Release History Table

Release	Description
19.1	Starting in Junos OS Release 19.1R1, segment routing-LDP border router can stitch segment routing traffic to LDP next-hop and vice versa.

RELATED DOCUMENTATION

[source-packet-routing](#) | [2983](#)
[mapping-server-entry](#) | [2949](#)

Configuring Miscellaneous LDP Properties

IN THIS SECTION

- [Configuring LDP to Use the IGP Route Metric](#) | [1116](#)
- [Preventing Addition of Ingress Routes to the inet.0 Routing Table](#) | [1116](#)
- [Multiple-Instance LDP and Carrier-of-Carriers VPNs](#) | [1116](#)
- [Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router](#) | [1117](#)
- [Enabling LDP over RSVP-Established LSPs](#) | [1117](#)
- [Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks](#) | [1118](#)
- [Configuring the TCP MD5 Signature for LDP Sessions](#) | [1118](#)
- [Configuring LDP Session Protection](#) | [1120](#)
- [Disabling SNMP Traps for LDP](#) | [1120](#)
- [Configuring LDP Synchronization with the IGP on LDP Links](#) | [1121](#)
- [Configuring LDP Synchronization with the IGP on the Router](#) | [1121](#)
- [Configuring the Label Withdrawal Timer](#) | [1121](#)
- [Ignoring the LDP Subnet Check](#) | [1122](#)

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

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.

NOTE: ACX Series routers do not support **[edit logical-systems]** hierarchy level.

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 User Guide*.

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

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

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

```
explicit-null;
```

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

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

For more information about labels, see [“MPLS Label Overview” on page 419](#) and [“MPLS Label Allocation” on page 420](#).

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 934](#)). 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](#) | 928

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]**

NOTE: ACX Series routers do not support **[edit logical-systems]** hierarchy level.

To enable LDP over RSVP LSPs, you also still need to complete the procedure in Section [“Enabling LDP over RSVP-Established LSPs”](#) on page 1117.

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.

Starting in Junos OS Release 16.1R1, support for Hashed Message Authentication Code (HMAC) and MD5 authentication for LDP sessions is extended from a per-session configuration to a subnet-match (that is, longest-prefix-match) configuration.

The support for subnet-match authentication provides flexibility in configuring authentication for automatically targeted LDP (TLDP) sessions, making the deployment of remote loop-free alternate (LFA) and FEC 129 pseudowires easy.

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

```
session-group prefix-length {
  authentication-key authentication-key;
}
```

Use the **session-group** 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. You must also configure the authentication algorithm by including the **authentication-algorithm algorithm** statement the **[edit protocols ldp]** hierarchy level.

```
[edit protocols ldp]
group group-name {
```



```

neighbor address {
  authentication-algorithm algorithm;
  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 MIB Explorer](#)..

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*]

NOTE: ACX Series routers do not support [edit logical-systems] hierarchy level.

Release History Table

Release	Description
16.1R1	Starting in Junos OS Release 16.1R1, support for Hashed Message Authentication Code (HMAC) and MD5 authentication for LDP sessions is extended from a per-session configuration to a subnet-match (that is, longest-prefix-match) configuration.

RELATED DOCUMENTATION

[LDP Introduction | 918](#)

[Junos OS LDP Protocol Implementation | 924](#)

[Minimum LDP Configuration | 933](#)

Configuring LDP LSP Traceroute

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

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

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

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


You can configure this statement at the following hierarchy levels:

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

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

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

Collecting LDP Statistics

IN THIS SECTION

- [LDP Statistics Output | 1125](#)
- [Disabling LDP Statistics on the Penultimate-Hop Router | 1126](#)
- [LDP Statistics Limitations | 1126](#)

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

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:

- **FEC**—FEC for which LDP traffic statistics are collected.
- **Type**—Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).
- **Packets**—Number of packets passed by the FEC since its LSP came up.
- **Bytes**—Number of bytes of data passed by the FEC since its LSP came up.
- **Shared**—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.
- **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.

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

IN THIS SECTION

- [Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels | 1127](#)
- [Tracing LDP Protocol Traffic Within FECs | 1128](#)
- [Examples: Tracing LDP Protocol Traffic | 1129](#)

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

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;
        }
    }
}

```


5

PART

MPLS Traffic Engineering

Understanding MPLS Traffic Engineering | **1133**

Configuring DiffServ-Aware Traffic Engineering to Achieve Service Level Guarantees
on an MPLS network | **1193**

Understanding MPLS Traffic Engineering

IN THIS CHAPTER

- [MPLS and Traffic Engineering | 1133](#)
- [MPLS Traffic Engineering and Signaling Protocols Overview | 1134](#)
- [Traffic Engineering Capabilities | 1136](#)
- [Components of Traffic Engineering | 1136](#)
- [Configuring Traffic Engineering for LSPs | 1137](#)
- [Enabling Interarea Traffic Engineering | 1140](#)
- [Enabling Inter-AS Traffic Engineering for LSPs | 1141](#)
- [Packet Forwarding Component | 1143](#)
- [Packet Forwarding Based on Label Swapping | 1144](#)
- [How a Packet Traverses an MPLS Backbone | 1144](#)
- [Information Distribution Component | 1145](#)
- [Path Selection Component | 1145](#)
- [Offline Path Planning and Analysis | 1146](#)
- [Signaling Component | 1147](#)
- [Flexible LSP Calculation and Configuration | 1147](#)
- [Link-State Distribution Using BGP Overview | 1148](#)
- [Example: Configuring Link State Distribution Using BGP | 1163](#)
- [Configuring Link State Distribution Using BGP | 1186](#)
- [Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages | 1189](#)

MPLS and Traffic Engineering

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

- Make more efficient use of expensive long-haul fibers.

- Control how traffic is rerouted in the face of single or multiple failures.
- Classify critical and regular traffic on a per-path basis.

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

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

This section discusses the following topics:

- [MPLS Label Overview on page 419](#)
- [MPLS Label Allocation on page 420](#)
- [Routers in an LSP on page 513](#)
- [How a Packet Travels Along an LSP on page 430](#)
- [Types of LSPs on page 431](#)
- [Scope of LSPs on page 431](#)
- [Constrained-Path LSP Computation on page 485](#)
- [Path Computation for LSPs on an Overloaded Router on page 503](#)
- [Computing Backup Paths for LSPs Using Fate Sharing on page 504](#)
- [Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts on page 504](#)
- [Advertising LSPs into IGP on page 508](#)

MPLS Traffic Engineering and Signaling Protocols Overview

Traffic engineering facilitates efficient and reliable network operations while simultaneously optimizing network resources and traffic performance. Traffic engineering provides the ability to move traffic flow away from the shortest path selected by the interior gateway protocol (IGP) to a potentially less congested physical path across a network. To support traffic engineering, besides source routing, the network must do the following:

- Compute a path at the source by taking into account all the constraints, such as bandwidth and administrative requirements.
- Distribute the information about network topology and link attributes throughout the network once the path is computed.
- Reserve network resources and modify link attributes.

When transit traffic is routed through an IP network, MPLS is often used to engineer its passage. Although the exact path through the transit network is of little importance to either the sender or the receiver of the traffic, network administrators often want to route traffic more efficiently between certain source and destination address pairs. By adding a short label with specific routing instructions to each packet, MPLS switches packets from router to router through the network rather than forwarding packets based on next-hop lookups. The resulting routes are called *label-switched paths (LSPs)*. LSPs control the passage of traffic through the network and speed traffic forwarding.

You can create LSPs manually, or through the use of signaling protocols. Signaling protocols are used within an MPLS environment to establish LSPs for traffic across a transit network. Junos OS supports two signaling protocols—LDP and the Resource Reservation Protocol (RSVP).

MPLS traffic engineering uses the following components:

- MPLS LSPs for packet forwarding
- IGP extensions for distributing information about the network topology and link attributes
- Constrained Shortest Path First (CSPF) for path computation and path selection
- RSVP extensions to establish the forwarding state along the path and to reserve resources along the path

Junos OS also supports traffic engineering across different OSPF regions.

RELATED DOCUMENTATION

[MPLS Applications User Guide for Routing Devices](#)

[Understanding the LDP Signaling Protocol | 919](#)

[Understanding the RSVP Signaling Protocol | 810](#)

[Understanding Point-to-Multipoint LSPs | 693](#)

Traffic Engineering Capabilities

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

Traffic engineering provides the capabilities to do the following:

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

Components of Traffic Engineering

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

- [Packet Forwarding Component on page 1143](#)
- [Information Distribution Component on page 1145](#)
- [Path Selection Component on page 1145](#)
- [Signaling Component on page 1147](#)

Configuring Traffic Engineering for LSPs

IN THIS SECTION

- Using LSPs for Both BGP and IGP Traffic Forwarding | 1137
- Using LSPs for Forwarding in Virtual Private Networks | 1138
- Using RSVP and LDP Routes for Forwarding but Not Route Selection | 1138
- Advertising the LSP Metric in Summary LSAs | 1139

When you configure an LSP, a host route (a 32-bit mask) is installed in the ingress router toward the egress router; the address of the host route is the destination address of the LSP. The **bgp** option for the **traffic-engineering** statement at the **[edit protocols mpls]** hierarchy level is enabled by default (you can also explicitly configure the **bgp** option), allowing only BGP to use LSPs in its route calculations. The other **traffic-engineering** statement options allow you to alter this behavior in the master routing instance. This functionality is not available for specific routing instances. Also, you can enable only one of the **traffic-engineering** statement options (**bgp**, **bgp-igp**, **bgp-igp-both-ribs**, or **mpls-forwarding**) at a time.

NOTE: Enabling or disabling any of the **traffic-engineering** statement options causes all the MPLS routes to be removed and then reinserted into the routing tables.

You can configure OSPF and traffic engineering to advertise the LSP metric in summary link-state advertisements (LSAs) as described in the section [“Advertising the LSP Metric in Summary LSAs” on page 1139](#).

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

Using LSPs for Both BGP and IGP Traffic Forwarding

You can configure BGP and the IGPs to use LSPs for forwarding traffic destined for egress routers by including the **bgp-igp** option for the **traffic-engineering** statement. The **bgp-igp** option causes all inet.3 routes to be moved to the inet.0 routing table.

On the ingress router, include **bgp-igp** option for the **traffic-engineering** statement:

```
traffic-engineering bgp-igp;
```

You can include this statement at the following hierarchy levels:

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

NOTE: The **bgp-igp** option for the **traffic-engineering** statement cannot be configured for VPN). VPNs require that routes be in the inet.3 routing table.

Using LSPs for Forwarding in Virtual Private Networks

VPNs require that routes remain in the inet.3 routing table to function properly. For VPNs, configure the **bgp-igp-both-ribs** option of the **traffic-engineering** statement to cause BGP and the IGP to use LSPs for forwarding traffic destined for egress routers. The **bgp-igp-both-ribs** option installs the ingress routes in both the inet.0 routing table (for IPv4 unicast routes) and the inet.3 routing table (for MPLS path information).

On the ingress router, include the **traffic-engineering bgp-igp-both-ribs** statement:

```
traffic-engineering bgp-igp-both-ribs;
```

You can include this statement at the following hierarchy levels:

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

When you use the **bgp-igp-both-ribs** statement, the routes from the inet.3 table get copied into the inet.0 table. The copied routes are LDP-signaled or RSVP-signaled, and are likely to have a lower preference than other routes in inet.0. Routes with a lower preference are more likely to be chosen as the active routes. This can be a problem because routing policies only act upon active routes. To prevent this problem, use the **mpls-forwarding** option instead.

Using RSVP and LDP Routes for Forwarding but Not Route Selection

If you configure the **bgp-igp** or **bgp-igp-both-ribs** options for the **traffic-engineering** statement, high-priority LSPs can supersede IGP routes in the inet.0 routing table. IGP routes might no longer be redistributed since they are no longer the active routes.

If you configure the **mpls-forwarding** option for the **traffic-engineering** statement, LSPs are used for forwarding but are excluded from route selection. These routes are added to both the inet.0 and inet.3 routing tables. LSPs in the inet.0 routing table are given a low preference when the active route is selected. However, LSPs in the inet.3 routing table are given a normal preference and are therefore used for selecting forwarding next hops.

When you activate the **mpls-forwarding** option, routes whose state is **ForwardingOnly** are preferred for forwarding even if their preference is lower than that of the currently active route. To examine the state of a route, execute a **show route detail** command.

To use LSPs for forwarding but exclude them from route selection, include the **mpls-forwarding** option for the **traffic-engineering** statement:

```
traffic-engineering mpls-forwarding;
```

You can include this statement at the following hierarchy levels:

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

When you configure the **mpls-forwarding** option, IGP shortcut routes are copied to the inet.0 routing table only.

Unlike the **bgp-igp-both-ribs** option, the **mpls-forwarding** option allows you to use the LDP-sigaled and RSVP-sigaled routes for forwarding, and keep the BGP and IGP routes active for routing purposes so that routing policies can act upon them.

For example, suppose a router is running BGP and it has a BGP route of 10.10.10.1/32 that it needs to send to another BGP speaker. If you use the **bgp-igp-both-ribs** option, and your router also has a label-switched-path (LSP) to 10.10.10.1, the MPLS route for 10.10.10.1 becomes active in the inet.0 routing table. This prevents your router from advertising the 10.10.10.1 route to the other BGP router. On the other hand, if you use the **mpls-forwarding** option instead of the **bgp-igp-both-ribs** option, the 10.10.10.1/32 BGP route is advertised to the other BGP speaker, and the LSP is still used to forward traffic to the 10.10.10.1 destination.

Advertising the LSP Metric in Summary LSAs

You can configure MPLS and OSPF to treat an LSP as a link. This configuration allows other routers in the network to use this LSP. To accomplish this goal, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

For MPLS, include the **traffic-engineering bgp-igp** and **label-switched-path** statements:

```
traffic-engineering bgp-igp;
label-switched-path lsp-name {
  to address;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls]

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

For OSPF, include the **lsp-metric-into-summary** statement:

```
lsp-metric-into-summary;
```

You can include this statement at the following hierarchy levels:

- [edit protocols ospf traffic-engineering shortcuts]
- [edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]

For more information about OSPF traffic engineering, see the *Junos OS Routing Protocols Library*.

Enabling Interarea Traffic Engineering

The Junos OS can signal a contiguous traffic-engineered LSP across multiple OSPF areas. The LSP signaling must be done using either nesting or contiguous signaling, as described in RFC 4206, *Label-Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*. However, contiguous signaling support is limited to just basic signaling. Reoptimization is not supported with contiguous signaling.

The following describes some of the interarea traffic engineering features:

- Interarea traffic engineering can be enabled when the loose-hop area border routers (ABRs) are configured on the ingress router using CSPF for the Explicit Route Object (ERO) calculation within an OSPF area. ERO expansion is completed on the ABRs.
- Interarea traffic engineering can be enabled when CSPF is enabled, but without ABRs specified in the LSP configuration on the ingress router (ABRs can be automatically designated).
- Differentiated Services (DiffServ) traffic engineering is supported as long as the class type mappings are uniform across multiple areas.

To enable interarea traffic engineering, include the **expand-loose-hop** statement in the configuration for each LSP transit router:

```
expand-loose-hop;
```

You can include this statement at the following hierarchy levels:

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

Enabling Inter-AS Traffic Engineering for LSPs

IN THIS SECTION

- [Inter-AS Traffic Engineering Requirements | 1141](#)
- [Inter-AS Traffic Engineering Limitations | 1142](#)
- [Configuring OSPF Passive TE Mode | 1143](#)

Generally, traffic engineering is possible for LSPs that meet the following conditions:

- Both ends of the LSP are in the same OSPF area or at the same IS-IS level.
- The two ends of the LSP are in different OSPF areas within the same autonomous system (AS). LSPs that end in different IS-IS levels are not supported.
- The two ends of an explicit-path LSP are in different OSPF ASs and the autonomous system border routers (ASBRs) are configured statically as the loose hops supported on the explicit-path LSP. For more information, see [“Configuring Explicit-Path LSPs” on page 686](#).

Without statically defined ASBRs on LSPs, traffic engineering is not possible between one routing domain, or AS, and another. However, when the ASs are under the control of single service provider, it is possible in some cases to have traffic engineered LSPs span the ASs and dynamically discover the OSPF ASBRs linking them (IS-IS is not supported with this feature).

Inter-AS traffic engineered LSPs are possible as long as certain network requirements are met, none of the limiting conditions apply, and OSPF passive mode is configured with EBGp. Details are provided in the following sections:

Inter-AS Traffic Engineering Requirements

The proper establishment and functioning of inter-AS traffic engineered LSPs depend on the following network requirements, all of which must be met:

- All ASs are under control of a single service provider.
- OSPF is used as the routing protocol within each AS, and EBGp is used as the routing protocol between the ASs.
- ASBR information is available inside each AS.
- EBGp routing information is distributed by OSPF, and an IBGP full mesh is in place within each AS.

- Transit LSPs are *not* configured on the inter-AS links, but *are* configured between entry and exit point ASBRs on each AS.
- The EBGP link between ASBRs in different ASs is a direct link and must be configured as a passive traffic engineering link under OSPF. The remote link address itself, not the loopback or any other link address, is used as the remote node identifier for this passive link. For more information about OSPF passive traffic engineering mode configuration, see [“Configuring OSPF Passive TE Mode” on page 1143](#).

In addition, the address used for the remote node of the OSPF passive traffic engineering link must be the same as the address used for the EBGP link. For more information about OSPF and BGP in general, see the *Junos OS Routing Protocols Library*.

Inter-AS Traffic Engineering Limitations

Only LSP hierarchical, or nested, signaling is supported for inter-AS traffic engineered LSPs. Only point-to-point LSPs are supported (there is no point-to-multipoint support).

In addition, the following limitations apply. Any one of these conditions is sufficient to render inter-AS traffic engineered LSPs impossible, even if the above requirements are met.

- The use of multihop BGP is not supported.
- The use of policers or topologies that prevent BGP routes from being known inside the AS is not supported.
- Multiple ASBRs on a LAN between EBGP peers are not supported. Only one ASBR on a LAN between EBGP peers is supported (others ASBRs can exist on the LAN, but cannot be advertised).
- Route reflectors or policies that hide ASBR information or prevent ASBR information from being distributed inside the ASs are not supported.
- Bidirectional LSPs are not supported (LSPs are unidirectional from the traffic engineering perspective).
- Topologies with both inter-AS and intra-AS paths to the same destination are not supported.

In addition, several features that are routine with all LSPs are not supported with inter-AS traffic engineering:

- Admin group link colors are not supported.
- Secondary standby is not supported.
- Reoptimization is not supported.
- Crankback on transit routers is not supported.
- Diverse path calculation is not supported.
- Graceful restart is not supported.

These lists of limitations or unsupported features with inter-AS traffic engineered LSPs are not exhaustive.

Configuring OSPF Passive TE Mode

Ordinarily, interior routing protocols such as OSPF are not run on links between ASs. However, for inter-AS traffic engineering to function properly, information about the inter-AS link, in particular, the address on the remote interface, must be made available inside the AS. This information is not normally included either in EBGp reachability messages or in OSPF routing advertisements.

To flood this link address information within the AS and make it available for traffic engineering calculations, you must configure OSPF passive mode for traffic engineering on each inter-AS interface. You must also supply the remote address for OSPF to distribute and include in the traffic engineering database.

To configure OSPF passive mode for traffic engineering on an inter-AS interface, include the **passive** statement for the link at the **[edit protocols ospf area *area-id* interface *interface-name*]** hierarchy level:

```
passive {
  traffic-engineering {
    remote-node-id ip-address; /* IP address at far end of inter-AS link */
  }
}
```

OSPF must be properly configured on the router. The following example configures the inter-AS link **so-1/1/0** to distribute traffic engineering information with OSPF within the AS. The remote IP address is **192.168.207.2**.

```
[edit protocols ospf area 0.0.0.0]
interface so-1/1/0 {
  unit 0 {
    passive {
      traffic-engineering {
        remote-node-id 192.168.207.2;
      }
    }
  }
}
```

Packet Forwarding Component

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

each direction. An LSP is created by the concatenation of one or more label-switched hops, allowing a packet to be forwarded from one router to another across the MPLS domain.

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

This section discusses the following topics:

- [Packet Forwarding Based on Label Swapping on page 1144](#)
- [How a Packet Traverses an MPLS Backbone on page 1144](#)

RELATED DOCUMENTATION

| [Components of Traffic Engineering | 1136](#)

Packet Forwarding Based on Label Swapping

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

How a Packet Traverses an MPLS Backbone

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

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

the ingress router that are destined to exit the MPLS domain at the same egress router are forwarded along the same LSP.

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

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

Information Distribution Component

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

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

RELATED DOCUMENTATION

[Components of Traffic Engineering](#) | 1136

Path Selection Component

After network link attributes and topology information are flooded by the IGP and placed in the traffic engineering database, each ingress router uses the traffic engineering database to calculate the paths for

its own set of LSPs across the routing domain. The path for each LSP can be represented by either a strict or loose explicit route. An explicit route is a preconfigured sequence of routers that should be part of the physical path of the LSP. If the ingress router specifies all the routers in the LSP, the LSP is said to be identified by a strict explicit route. If the ingress router specifies only some of the routers in the LSP, the LSP is described as a loose explicit route. Support for strict and loose explicit routes allows the path selection process to be given broad latitude whenever possible, but to be constrained when necessary.

The ingress router determines the physical path for each LSP by applying a Constrained Shortest Path First (CSPF) algorithm to the information in the traffic engineering database. CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. Input into the CSPF algorithm includes:

- Topology link-state information learned from the IGP and maintained in the traffic engineering database
- Attributes associated with the state of network resources (such as total link bandwidth, reserved link bandwidth, available link bandwidth, and link color) that are carried by IGP extensions and stored in the traffic engineering database
- Administrative attributes required to support traffic traversing the proposed LSP (such as bandwidth requirements, maximum hop count, and administrative policy requirements) that are obtained from user configuration

As CSPF considers each candidate node and link for a new LSP, it either accepts or rejects a specific path component based on resource availability or whether selecting the component violates user policy constraints. The output of the CSPF calculation is an explicit route consisting of a sequence of router addresses that provides the shortest path through the network that meets the constraints. This explicit route is then passed to the signaling component, which establishes the forwarding state in the routers along the LSP.

RELATED DOCUMENTATION

| [Components of Traffic Engineering](#) | 1136

Offline Path Planning and Analysis

Despite the reduced management effort resulting from online path calculation, an offline planning and analysis tool is still required to optimize traffic engineering globally. Online calculation takes resource constraints into account and calculates one LSP at a time. The challenge with this approach is that it is not deterministic. The order in which LSPs are calculated plays a critical role in determining each LSP's physical path across the network. LSPs that are calculated early in the process have more resources available to them than LSPs calculated later in the process because previously calculated LSPs consume network

resources. If the order in which the LSPs are calculated is changed, the resulting set of physical paths for the LSPs also can change.

An offline planning and analysis tool simultaneously examines each link's resource constraints and the requirements of each LSP. Although the offline approach can take several hours to complete, it performs global calculations, compares the results of each calculation, and then selects the best solution for the network as a whole. The output of the offline calculation is a set of LSPs that optimizes utilization of network resources. After the offline calculation is completed, the LSPs can be established in any order because each is installed according to the rules for the globally optimized solution.

Signaling Component

An LSP is not known to be workable until it is actually established by the signaling component. The signaling component, which is responsible for establishing LSP state and distributing labels, relies on a number of extensions to RSVP:

- The Explicit Route object allows an RSVP path message to traverse an explicit sequence of routers that is independent of conventional shortest-path IP routing. The explicit route can be either strict or loose.
- The Label Request object permits the RSVP path message to request that intermediate routers provide a label binding for the LSP that it is establishing.
- The Label object allows RSVP to support the distribution of labels without changing its existing mechanisms. Because the RSVP Resv message follows the reverse path of the RSVP path message, the Label object supports the distribution of labels from downstream nodes to upstream nodes.

RELATED DOCUMENTATION

[Components of Traffic Engineering](#) | 1136

Flexible LSP Calculation and Configuration

Traffic engineering involves mapping traffic flow onto a physical topology. You can determine the paths online using constraint-based routing. Regardless of how the physical path is calculated, the forwarding state is installed across the network through RSVP.

The Junos OS supports the following ways to route and configure an LSP:

- You can calculate the full path for the LSP offline and individually configure each router in the LSP with the necessary static forwarding state. This is analogous to the way some Internet service providers (ISPs) configure their IP-over-ATM cores.
- You can calculate the full path for the LSP offline and statically configure the ingress router with the full path. The ingress router then uses RSVP as a dynamic signaling protocol to install a forwarding state in each router along the LSP.
- You can rely on constraint-based routing to perform dynamic online LSP calculation. You configure the constraints for each LSP; then the network itself determines the path that best meets those constraints. Specifically, the ingress router calculates the entire LSP based on the constraints and then initiates signaling across the network.
- You can calculate a partial path for an LSP offline and statically configure the ingress router with a subset of the routers in the path; then you can permit online calculation to determine the complete path.

For example, consider a topology that includes two east-west paths across the United States: one in the north through Chicago and one in the south through Dallas. If you want to establish an LSP between a router in New York and one in San Francisco, you can configure the partial path for the LSP to include a single loose-routed hop of a router in Dallas. The result is an LSP routed along the southern path. The ingress router uses CSPF to compute the complete path and RSVP to install the forwarding state along the LSP.

- You can configure the ingress router with no constraints whatsoever. In this case, normal IGP shortest-path routing is used to determine the path of the LSP. This configuration does not provide any value in terms of traffic engineering. However, it is easy and might be useful in situations when services such as virtual private networks (VPNs) are needed.

In all these cases, you can specify any number of LSPs as backups for the primary LSP, thus allowing you to combine more than one configuration approach. For example, you might explicitly compute the primary path offline, set the secondary path to be constraint-based, and have the tertiary path be unconstrained. If a circuit on which the primary LSP is routed fails, the ingress router notices the outage from error notifications received from a downstream router or by the expiration of RSVP soft-state information. Then the router dynamically forwards traffic to a hot-standby LSP or calls on RSVP to create a forwarding state for a new backup LSP.

Link-State Distribution Using BGP Overview

IN THIS SECTION

- [Role of an Interior Gateway Protocol | 1149](#)
- [Limitations of an Interior Gateway Protocol | 1150](#)
- [Need for Spanning Link-State Distribution | 1150](#)

- [Using BGP as a Solution | 1150](#)
- [Supported and Unsupported Features | 1156](#)
- [BGP Link-State Extensions for Source Packet Routing in Networking \(SPRING\) | 1157](#)
- [Verifying NLRI Node Learned Through BGP with OSPF as IGP | 1160](#)
- [Verifying the Prefix NLRI Learned Through BGP with OSPF as IGP | 1161](#)

Role of an Interior Gateway Protocol

An interior gateway protocol (IGP) is a type of protocol used for exchanging routing information between devices within an autonomous system (AS). Based on the method of computing the best path to a destination, the IGPs are divided into two categories:

- **Link-state protocols**—Advertise information about the network topology (directly connected links and the state of those links) to all routers using multicast addresses and triggered routing updates until all the routers running the link-state protocol have identical information about the internetwork. The best path to a destination is calculated based on constraints such as maximum delay, minimum available bandwidth, and resource class affinity.

OSPF and IS-IS are examples of link-state protocols.

- **Distance vector protocols**—Advertise complete routing table information to directly connected neighbors using a broadcast address. The best path is calculated based on the number of hops to the destination network.

RIP is an example of a distance vector protocol.

As the name implies, the role of an IGP is to provide routing connectivity within or internal to a given routing domain. A routing domain is a set of routers under common administrative control that share a common routing protocol. An AS can consist of multiple routing domains, where IGP functions to advertise and learn network prefixes (routes) from neighboring routers to build a route table that ultimately contains entries for all sources advertising reachability for a given prefix. IGP executes a route selection algorithm to select the best path between the local router and each destination, and provides full connectivity among the routers making up a routing domain.

In addition to advertising internal network reachability, IGPs are often used to advertise routing information that is external to that IGP's routing domain through a process known as route redistribution. Route redistribution is the process of exchanging routing information among distinct routing protocols to tie multiple routing domains together when intra-AS connectivity is desired.

Limitations of an Interior Gateway Protocol

While each individual IGP has its own advantages and limitations, the biggest limitations of IGP in general are performance and scalability.

IGPs are designed to handle the task of acquiring and distributing network topology information for traffic engineering purposes. While this model has served well, IGPs have inherent scaling limitations when it comes to distributing large databases. IGPs can autodetect neighbors, with which they acquire intra-area network topology information. However, the link-state database or a traffic engineering database has the scope of a single area or AS, thereby limiting applications, such as end-to-end traffic engineering, the benefit of having external visibility to make better decisions.

For label-switched networks, such as MPLS and Generalized MPLS (GMPLS), most existing traffic engineering solutions work in a single routing domain. These solutions do not work when a route from the ingress node to the egress node leaves the routing area or AS of the ingress node. In such cases, the path computation problem becomes complicated because of the unavailability of the complete routing information throughout the network. This is because service providers usually choose not to leak routing information beyond the routing area or AS for scalability constraints and confidentiality concerns.

Need for Spanning Link-State Distribution

One of the limitations of IGP is its inability to span link-state distribution outside a single area or AS. However, spanning link-state information acquired by an IGP across multiple areas or ASs has the following needs:

- LSP path computation—This information is used to compute the path for MPLS LSPs across multiple routing domains, for example an inter-area TE LSP.
- External path computing entities—External path computing entities, such as Application Layer Traffic Optimization (ALTO) and Path Computation Elements (PCE), perform path computations based on the network topology and current state of connections within the network, including traffic engineering information. This information is typically distributed by IGPs within the network.

However, because the external path computing entities cannot extract this information from the IGPs, they perform network monitoring to optimize network services.

Using BGP as a Solution

IN THIS SECTION

- [Overview | 1151](#)
- [Implementation | 1151](#)

Overview

To meet the needs for spanning link-state distribution across multiple domains, an exterior gateway protocol (EGP) is required to collect link-state and traffic engineering information from an IGP area, share it with external component, and use it for computing paths for interdomain MPLS LSPs.

BGP is a standardized EGP designed to exchange routing and reachability information between autonomous systems (ASs). BGP is a proven protocol that has better scaling properties because it can distribute millions of entries (for example, VPN prefixes) in a scalable fashion. BGP is the only routing protocol in use today that is suited to carry all of the routes in the Internet. This is largely because BGP runs on top of TCP and can make use of TCP flow control. In contrast, the internal gateway protocols (IGPs) do not have flow control. When IGPs have too much route information, they begin to churn. When BGP has a neighboring speaker that is sending information too quickly, BGP can throttle down the neighbor by delaying TCP acknowledgments.

Another benefit of BGP is that it uses type, length, value (TLV) tuples and network layer reachability information (NLRI) that provide seemingly endless extensibility without the need for the underlying protocol to be altered.

The distribution of link-state information across domains is regulated using policies to protect the interests of the service provider. This requires a control over the topology distribution using policies. BGP with its implemented policy framework serves well in the interdomain route distribution. In Junos OS, BGP is completely policy driven. The operator must explicitly configure neighbors to peer with and explicitly accept routes into BGP. Furthermore, routing policy is used to filter and modify routing information. Thus, routing policies provide complete administrative control over the routing tables.

Although, within an AS, both IGP-TE and BGP-TE provide the same set of information, BGP-TE has better scaling characteristics that are inherited from the standard BGP protocol. This makes BGP-TE a more scalable choice for acquiring multi-area/multi-AS topology information.

By using BGP as a solution, the IGP-acquired information is used for distribution into BGP. The ISPs can selectively expose this information with other ISPs, service providers, and content distribution networks (CDNs) through normal BGP peering. This allows for aggregation of the IGP-acquired information across multiple areas and ASs, such that an external path computing entity can access the information by passively listening to a route reflector.

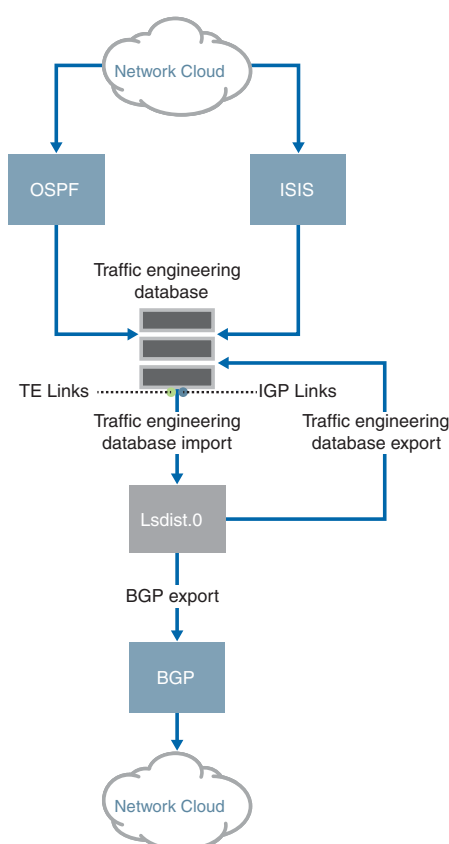
Implementation

IN THIS SECTION

- [Traffic Engineering Database Import | 1152](#)
- [Traffic Engineering Database Export | 1153](#)
- [Assigning Credibility Values | 1153](#)
- [Cross-Credibility Path Computation | 1154](#)
- [BGP-TE NLRIs and TLVs | 1154](#)

In Junos OS, the IGP's install topology information into a database called the traffic engineering database. The traffic engineering database contains the aggregated topology information. To install IGP topology information into traffic engineering database, use the **set igp-topology** configuration statement at the **[edit protocols isis traffic-engineering]** and **[edit protocols ospf traffic-engineering]** hierarchy levels. The mechanism to distribute link-state information using BGP includes the process of advertising the traffic engineering database into BGP-TE (import), and installing entries from BGP-TE into the traffic engineering database (export).

Figure 96: Junos OS Implementation of BGP Link-State Distribution



Traffic Engineering Database Import

To advertise the traffic engineering database into BGP-TE, the link and node entries in the traffic engineering database are converted in the form of routes. These converted routes are then installed by the traffic engineering database on behalf of the corresponding IGP, into a user-visible routing table called **Lsdist.0**, on conditions subject to route policies. The procedure of leaking entries from the traffic engineering database into **Lsdist.0** is called traffic engineering database import as illustrated in [Figure 96 on page 1152](#).

There are policies to govern the traffic engineering database import process. By default, no entries are leaked from the traffic engineering database into the **Lsdist.0** table.

Starting in Junos OS Release 17.4R1, the traffic engineering database installs interior gateway protocol (IGP) topology information in addition to RSVP-TE topology information in the **Isdist.0** routing table as illustrated in [Figure 96 on page 1152](#). Prior to Junos OS Release 17.4R1, the traffic engineering database only exported RSVP-TE topology information. Now you can monitor both IGP and traffic engineering topology information. The BGP-LS reads IGP entries from **Isdist.0** and advertises these entries to the BGP peers. To import IGP topology information into BGP-LS from **Isdist.0**, use the **set bgp-ls** configuration statement at the **[edit protocols mpls traffic-engineering database import igp-topology]** hierarchy level.

Traffic Engineering Database Export

BGP can be configured to export or advertise routes from the **Isdist.0** table, subject to policy. This is common for any kind of route origination in BGP. In order to advertise BGP-TE into the traffic engineering database, BGP needs to be configured with the BGP-TE address family, and an export policy that selects routes for redistribution into BGP.

BGP then propagates these routes like any other NLRI. BGP peers that have the BGP-TE family configured and negotiated receive BGP-TE NLRIs. BGP stores the received BGP-TE NLRIs in the form of routes in the **Isdist.0** table, which is the same table that stores locally originated BGP-TE routes. The BGP-installed routes in **Isdist.0** are then distributed to other peers like any other route. Thus, the standard route selection procedure applies to BGP-TE NLRIs received from multiple speakers.

To achieve interdomain TE, the routes in **Isdist.0** are leaked into the traffic engineering database through a policy. This process is called traffic engineering database export as illustrated in [Figure 96 on page 1152](#).

There are policies to govern the traffic engineering database export process. By default, no entries are leaked from the **Isdist.0** table into the traffic engineering database.

NOTE: For SDN applications, such as PCE and ALTO, the BGP-TE advertised information cannot leak into the traffic engineering database of a router. In such cases, an external server that peers with the routers using BGP-TE is used to move topology information up into the sky/orchestration system that spans the network. These external servers can be deemed as BGP-TE consumers, where they receive BGP-TE routes, but do not advertise them.

Assigning Credibility Values

Once the entries are installed in the traffic engineering database, the BGP-TE learned information is made available for CSPF path computation. The traffic engineering database uses a protocol preference scheme that is based on credibility values. A protocol with a higher credibility value is preferred over a protocol with a lower credibility value. BGP-TE has the capability to advertise information learned from multiple protocols at the same time, and so in addition to the IGP-installed entries in the traffic engineering database, there can be BGP-TE installed entries that correspond to more than one protocol. The traffic engineering database export component creates a traffic engineering database protocol and credibility level for each protocol that BGP-TE supports. These credibility values are configurable in the CLI.

The credibility order for the BGP-TE protocols is as follows:

- Unknown—80
- OSPF—81
- ISIS Level 1—82
- ISIS Level 2—83
- Static—84
- Direct—85

Cross-Credibility Path Computation

After you assign credibility values, each credibility level is treated as an individual plane. The Constrained Shorted Path First algorithm starts with the highest assigned credibility to the lowest, finding a path within that credibility level.

With BGP-TE, it is essential to compute paths across credibility levels to compute inter-AS paths. For example, different credibility settings are seen on a device from area 0 that computes the path through area 1, because area 0 entries are installed by OSPF, and area 1 entries are installed by BGP-TE.

To enable path computation across credibility levels, include the **cross-credibility-cspf** statement at the **edit protocols mpls**, **[edit protocols mpls label-switched-path lsp-name]**, and **[edit protocols rsvp]** hierarchy levels. At the **[edit protocols rsvp]** hierarchy level, enabling **cross-credibility-cspf** impacts bypass LSPs and loose hop expansion in transit.

Configuring **cross-credibility-cspf** enables path computation across credibility levels using the Constrained Shortest Path First algorithm, wherein the constraint is not performed on a credibility-by-credibility basis, but as a single constraint ignoring the assigned credibility values.

BGP-TE NLRIs and TLVs

Like other BGP routes, BGP-TE NLRIs can also be distributed through a route reflector that speaks BGP-TE NLRI. Junos OS implements the route reflection support for the BGP-TE family.

The following is a list of supported NLRIs:

- Link NLRI
- Node NLRI
- IPv4 Prefix NLRI (receive and propagate)
- IPv6 Prefix NLRI (receive and propagate)

NOTE: Junos OS does not provide support for the route-distinguisher form of the above NLRIs.

The following is a list of supported fields in link and node NLRIs:

- Protocol-ID—NLRI originates with the following protocol values:
 - ISIS-L1
 - ISIS-L2
 - OSPF
- Identifier—This value is configurable. By default, the identifier value is set to 0.
- Local/Remote node descriptor—These include:
 - Autonomous system
 - BGP-LS Identifier—This value is configurable. By default, the BGP-LS identifier value is set to 0
 - Area-ID
 - IGP router-ID
- Link descriptors (Only for link NLRI)—This includes:
 - Link Local/Remote Identifiers
 - IPv4 interface address
 - IPv4 neighbor address
 - IPv6 neighbor/interface address—The IPv6 neighbor and interface addresses are not originated, but only stored and propagated when received.
 - Multi-topology ID—This value is not originated, but stored and propagated when received.

The following is a list of supported LINK_STATE attribute TLVs:

- Link attributes:
 - Administrative group
 - Max link bandwidth
 - Max reservable bandwidth
 - Unreserved bandwidth
 - TE default metric
 - SRLG
- The following TLVs, which are not originated, but only stored and propagated when received:
 - Opaque link attributes
 - MPLS protocol mask
 - Metric

- Link protection type
- Link name attribute
- Node attributes:
 - IPv4 Router-ID
 - Node flag bits—Only the overload bit is set.
 - The following TLVs, which are not originated, but only stored and propagated when received:
 - Multi-topology
 - OSPF-specific node properties
 - Opaque node properties
 - Node name
 - IS-IS area identifier
 - IPv6 Router-ID
 - Prefix attributes—These TLVs are stored and propagated like any other unknown TLVs.

Supported and Unsupported Features

Junos OS supports the following features with link-state distribution using BGP:

- Advertisement of multiprotocol assured forwarding capability
- Transmission and reception of node and link-state BGP and BGP-TE NLRIs
- Nonstop active routing for BGP-TE NLRIs
- Policies

Junos OS does **not** support the following functionality for link-state distribution using BGP:

- Aggregated topologies, links, or nodes
- Route distinguisher support for BGP-TE NLRIs
- Multi-topology identifiers
- Multi-instance identifiers (excluding the default instance ID 0)
- Advertisement of the link and node area TLV
- Advertisement of MPLS signaling protocols
- Importing node and link information with overlapping address

BGP Link-State Extensions for Source Packet Routing in Networking (SPRING)

IN THIS SECTION

- [Source Packet Routing in Networking \(SPRING\) | 1157](#)
- [Flow of BGP Link-State SPRING Data | 1157](#)
- [Supported BGP Link-State Attributes and TLVs, and Unsupported Features for BGP Link-State with SPRING | 1159](#)

Starting in Junos OS Release 17.2R1, the BGP link-state address family is extended to distribute the source packet routing in networking (SPRING) topology information to software-defined networking (SDN) controllers. BGP typically learns the link-state information from IGP and distributes it to BGP peers. Besides BGP, the SDN controller can get link-state information directly from IGP if the controller is a part of an IGP domain. However, BGP link-state distribution provides a scalable mechanism to export the topology information. BGP link-state extensions for SPRING is supported on interdomain networks.

Source Packet Routing in Networking (SPRING)

SPRING is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to decide the actual path it must take. SPRING engages IGPs, such as IS-IS and OSPF, for advertising network segments. Network segments can represent any instruction, topological or service-based. Within IGP topologies, IGP segments are advertised by the link-state routing protocols. There are two types of IGP segments:

Adjacency segment—A one-hop path over a specific adjacency between two nodes in the IGP

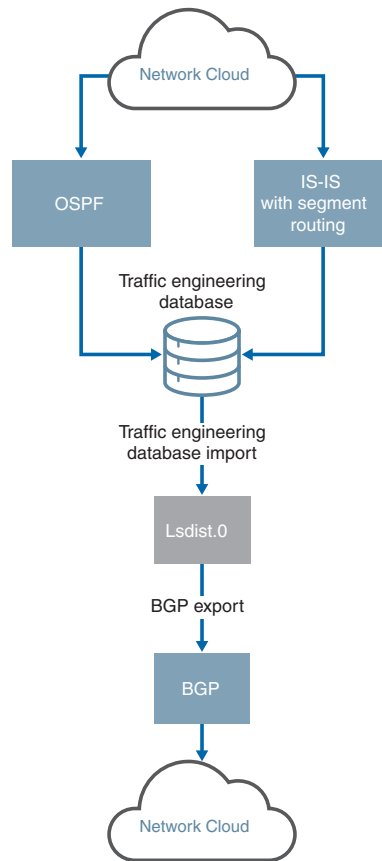
Prefix segment—A multi-hop, equal-cost, multipath-aware shortest-path to a prefix, as per the state of the IGP topology

When SPRING is enabled in a BGP network, BGP link-state address family learns the SPRING information from the IGP link-state routing protocols and advertises segments in the form of segment identifiers (SIDs). BGP link-state address family has been extended to carry SIDs and other SPRING-related information to BGP peers. The route reflector can steer a packet through a desired set of nodes and links by prepending the packet with an appropriate combination of tunnels. This feature allows BGP link-state address family to also advertise the SPRING information to BGP peers.

Flow of BGP Link-State SPRING Data

[Figure 97 on page 1158](#) depicts the data flow of BGP link-state SPRING data that IS-IS pushes to the traffic engineering database.

Figure 97: BGP Link-State Source Packet Routing in Networking (SPRING)



g043722

- IGP pushes the SPRING attributes to the traffic engineering database.
- SPRING capabilities and algorithm information are carried forward as node attributes into the traffic engineering database.
- Adjacent SID and LAN adjacent SID information are carried as link attributes.
- Prefix SID or node-SID information is carried as prefix attributes.
- A new set or a change to existing attributes triggers IGP updates to the traffic engineering database with new data.
- RSVP is a prerequisite for link attributes.



CAUTION: If traffic engineering is disabled at the IGP level, none of the attributes are pushed to the traffic engineering database.

- All parameters in the BGP traffic engineering NLRI, including the link, node, and prefix descriptors are derived from entries in the traffic engineering database.

- The traffic engineering database imports route entries into the **Isdist.0** routing table from IGP subject to policy.
- The default policy of BGP is to export routes, which are known to BGP only. You configure an export policy for non-BGP routes in the **Isdis.0** routing table. This policy advertises an entry learned from the traffic engineering database.

Supported BGP Link-State Attributes and TLVs, and Unsupported Features for BGP Link-State with SPRING

BGP link-state with SPRING supports the following attributes and type, length, and values (TLVs) that are originated, received, and propagated in the network:

Node attributes

- Segment routing Capabilities
- Segment routing Algorithm

Link attributes

- Adjacent-SID
- LAN Adjacent-SID

Prefix descriptors

- IP reachability information

Prefix attributes

- Prefix SID

The following list supports TLVs that are not originated, but only received and propagated in the network:

Prefix descriptors

- Multitopology ID
- OSPF route type

Prefix attributes

- Range
- Binding SID

Junos OS does not support the following features with BGP link-state with SPRING extensions:

- IPv6 prefix origination
- Multitopology identifiers
- Traffic engineering database export for SPRING parameters

- New TLVs with tcpdump (existing TLVs are also not supported).
- SPRING over IPv6

Verifying NLRI Node Learned Through BGP with OSPF as IGP

The following is a sample output to verify the NLRI node learned through BGP with OSPF as the IGP:

Purpose

Verify the lsdist.0 routing table entries.

Action

From operational mode, run the **show route table lsdist.0** command.

user@host> **show route table lsdist.0 te-node-ip 7.7.7.7 extensive**

```
lsdist.0: 216 destinations, 216 routes (216 active, 0 holddown, 0 hidden)
NODE { AS:100 Area:0.0.0.1 IPv4:7.7.7.7 OSPF:0 }/1536 (1 entry, 1 announced)
TSI:
LINK-STATE attribute handle 0x61d5da0
    *BGP      Preference: 170/-101
              Next hop type: Indirect, Next hop index: 0
              Address: 0x61b07cc
              Next-hop reference count: 216
              Source: 2.2.2.2
              Protocol next hop: 2.2.2.2
              Indirect next hop: 0x2 no-forward INH Session ID: 0x0
              State:<Active Int Ext>
              Local AS:   100 Peer AS:   100
              Age: 30:22      Metric2: 2
              Validation State: unverified
              Task: BGP_100.2.2.2.2
              Announcement bits (1): 0-TED Export
              AS path: I
              Accepted
              Area border router: No
              External router: No
              Attached: No
              Overload: No
              SPRING-Capabilities:
                - SRGB block [Start: 900000, Range: 90000, Flags: 0x00]
              SPRING-Algorithms:
                - Algo: 0
              Localpref: 100
```



```

Router ID: 2.2.2.2
Indirect next hops: 1
    Protocol next hop: 2.2.2.2 Metric: 2
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 11.1.1.2 via et-0/0/0.1 weight 0x1
        Session Id: 0x143
        2.2.2.2/32 Originating RIB: inet.0
        Metric: 2      Node path count: 1
        Forwarding nexthops: 1
            Nexthop: 11.1.1.2 via et-0/0/0.1
            Session Id: 143

```

Meaning

The routes are appearing in the lsdist.0 routing table.

Verifying the Prefix NLRI Learned Through BGP with OSPF as IGP

The following is a sample output to verify the prefix NLRI learned through BGP with OSPF as the IGP:

Purpose

Verify the lsdist.0 routing table entries.

Action

From operational mode, run the **show route table lsdist.0** command.

user@host> **show route table lsdist.0 te-ipv4-prefix-node-ip 7.7.7.7 extensive**

```

lsdist.0: 216 destinations, 216 routes (216 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:100 Area:0.0.0.1 IPv4:7.7.7.7 } { IPv4:7.7.7.7/32 } OSPF:0
}/1536 (1 entry, 0 announced)
    *BGP      Preference: 170/-101
        Next hop type: Indirect, Next hop index: 0
        Address: 0x61b07cc
        Next-hop reference count: 216
        Source: 2.2.2.2
        Protocol next hop: 2.2.2.2
        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
        State: <Active Int Ext>
        Local AS: 100 Peer AS: 100
        Age: 30:51      Metric2: 2

```



```
Validation State: unverified
Task: BGP_100.2.2.2.2
AS path: I
Accepted
Prefix Flags: 0x00, Prefix SID: 1007, Flags: 0x50, Algo: 0
Localpref: 100
Router ID: 2.2.2.2
Indirect next hops: 1
    Protocol next hop: 2.2.2.2 Metric: 2
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 11.1.1.2 via et-0/0/0.1 weight 0x1
        Session Id: 0x143
        2.2.2.2/32 Originating RIB: inet.0
        Metric: 2      Node path count: 1
        Forwarding nexthops: 1
            Nexthop: 11.1.1.2 via et-0/0/0.1
            Session Id: 143
```

Meaning

The routes are appearing in the lsdist.0 routing table.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, the traffic engineering database installs interior gateway protocol (IGP) topology information in addition to RSVP-TE topology information in the lsdist.0 routing table
17.2R1	Starting in Junos OS Release 17.2R1, the BGP link-state address family is extended to distribute the source packet routing in networking (SPRING) topology information to software-defined networking (SDN) controllers.

RELATED DOCUMENTATION

Example: Configuring Link State Distribution Using BGP 1163
<i>ipv4-prefix</i>

Example: Configuring Link State Distribution Using BGP

IN THIS SECTION

- [Requirements | 1163](#)
- [Overview | 1164](#)
- [Configuration | 1164](#)
- [Verification | 1178](#)

This example shows how to configure BGP to carry link-state information across multiple domains, which is used for computing paths for MPLS LSPs spanning multiple domains, such as inter-area TE LSP, and providing a scalable and policy-controlled means for external path computing entities, such as ALTO and PCE, to acquire network topology.

Requirements

This example uses the following hardware and software components:

- Four routers that can be a combination of M Series, MX Series, or T Series routers
- Junos OS Release 14.2 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
 - RSVP
 - MPLS
 - BGP
 - IS-IS
 - OSPF


Overview

Starting with Junos OS Release 14.2, a new mechanism to distribute topology information across multiple areas and autonomous systems (ASs) is introduced by extending the BGP protocol to carry link -state information, which was initially acquired using IGP. The IGP protocols have scaling limitations when it comes to distributing large databases. BGP is not only a more scalable vehicle for carrying multi-area and multi-AS topology information, but also provides the policy controls that can be useful for multi-AS topology distribution. The BGP link-state topology information is used for computing paths for MPLS label-switched paths (LSPs) spanning multiple domains, such as inter-area TE LSP, and providing a scalable and policy-controlled means for external path computing entities, such as ALTO and PCE, to acquire network topology.

Starting with Junos OS Release 17.1R1, link state distribution using BGP is supported on QFX10000 switches.

Topology

Figure 98: Link-State Distribution Using BGP



In [Figure 98 on page 1164](#), Routers R0 and R1 and Routers R2 and R3 belong to different autonomous systems. Routers R0 and R1 run OSPF, and Routers R2 and R3 run IS-IS.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

R0


```

set interfaces ge-0/0/0 unit 0 family inet address 8.31.1.101/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.105.137/32
set routing-options router-id 10.255.105.137
set routing-options autonomous-system 65533
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering database export policy accept-all
set protocols mpls cross-credibility-cspf
set protocols mpls label-switched-path to-R3-inter-as to 10.255.105.135
set protocols mpls label-switched-path to-R3-inter-as bandwidth 40m
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.105.137
set protocols bgp group ibgp family traffic-engineering unicast
set protocols bgp group ibgp neighbor 10.255.105.141
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set policy-options policy-statement accept-all from family traffic-engineering
set policy-options policy-statement accept-all then accept

```

R1

```

set interfaces ge-0/0/0 unit 0 family inet address 8.31.1.103/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 8.42.1.102/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.105.141/32
set routing-options router-id 10.255.105.141
set routing-options autonomous-system 65533
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal

```



```

set protocols bgp group ibgp local-address 10.255.105.141
set protocols bgp group ibgp family traffic-engineering unicast
set protocols bgp group ibgp export nlri2bgp
set protocols bgp group ibgp neighbor 10.255.105.137
set protocols bgp group ebgp type external
set protocols bgp group ebgp family traffic-engineering unicast
set protocols bgp group ebgp neighbor 8.42.1.104 local-address 8.42.1.102
set protocols bgp group ebgp neighbor 8.42.1.104 peer-as 65534
set protocols isis interface ge-0/0/1.0 passive remote-node-iso 0102.5502.4211
set protocols isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.104
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-id
    8.42.1.104
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-router-id
    10.255.105.139
set policy-options policy-statement accept-all from family traffic-engineering
set policy-options policy-statement accept-all then accept
set policy-options policy-statement nlri2bgp term 1 from family traffic-engineering
set policy-options policy-statement nlri2bgp term 1 then accept

```

R2

```

set interfaces ge-0/0/0 unit 0 family inet address 8.64.1.104/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 8.42.1.104/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.105.139/32
set interfaces lo0 unit 0 family iso
set routing-options router-id 10.255.105.139
set routing-options autonomous-system 65534
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering database import policy ted2nlri
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ebgp type external

```



```

set protocols bgp group ebgp family traffic-engineering unicast
set protocols bgp group ebgp export nlri2bgp
set protocols bgp group ebgp peer-as 65533
set protocols bgp group ebgp neighbor 8.42.1.102
set protocols isis level 1 disable
set protocols isis interface ge-0/0/0.0
set protocols isis interface ge-0/0/1.0 passive remote-node-iso 0102.5501.8181
set protocols isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.102
set protocols isis interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-id
    8.42.1.102
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-router-id
    10.255.105.141
set policy-options policy-statement accept-all from family traffic-engineering
set policy-options policy-statement accept-all then accept
set policy-options policy-statement nlri2bgp term 1 from family traffic-engineering
set policy-options policy-statement nlri2bgp term 1 then accept
set policy-options policy-statement ted2nlri term 1 from protocol isis
set policy-options policy-statement ted2nlri term 1 from protocol ospf
set policy-options policy-statement ted2nlri term 1 then accept
set policy-options policy-statement ted2nlri term 2 then reject

```

R3

```

set interfaces ge-0/0/0 unit 0 family inet address 8.64.1.106/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.105.135/32
set interfaces lo0 unit 0 family iso
set routing-options router-id 10.255.105.135
set routing-options autonomous-system 65534
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering database export policy accept-all
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.105.135
set protocols bgp group ibgp family traffic-engineering unicast

```



```

set protocols bgp group ibgp neighbor 10.255.105.139
set protocols isis interface ge-0/0/0.0 level 1 disable
set protocols isis interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set policy-options policy-statement accept-all from family traffic-engineering
set policy-options policy-statement accept-all then accept

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure Router R1:

1. Configure the Router R1 interfaces.

```

[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet address 8.31.1.103/24
user@R1# set ge-0/0/0 unit 0 family iso
user@R1# set ge-0/0/0 unit 0 family mpls
user@R1# set ge-0/0/1 unit 0 family inet address 8.42.1.102/24
user@R1# set ge-0/0/1 unit 0 family iso
user@R1# set ge-0/0/1 unit 0 family mpls
user@R1# set lo0 unit 0 family inet address 10.255.105.141/32

```

2. Configure the router ID and autonomous system of Router R1.

```

[edit routing-options]
user@R1# set router-id 10.255.105.141
user@R1# set autonomous-system 65533

```

3. Enable RSVP on all the interfaces of Router R1 (excluding the management interface).

```

[edit protocols]
user@R1# set rsvp interface all
user@R1# set rsvp interface fxp0.0 disable

```

4. Enable MPLS on all the interfaces of Router R1 (excluding the management interface).


```
[edit protocols]
user@R1# set mpls interface all
user@R1# set mpls interface fxp0.0 disable
```

5. Configure the BGP group for Router R1 to peer with Router R0, and assign the local address and neighbor address.

```
[edit protocols]
user@R1# set bgp group ibgp type internal
user@R1# set bgp group ibgp local-address 10.255.105.141
user@R1# set bgp group ibgp neighbor 10.255.105.137
```

6. Include the BGP-TE signaling network layer reachability information (NLRI) to the ibgp BGP group.

```
[edit protocols]
user@R1# set bgp group ibgp family traffic-engineering unicast
```

7. Enable export of policy nlri2bgp on Router R1.

```
[edit protocols]
user@R1# set bgp group ibgp export nlri2bgp
```

8. Configure the BGP group for Router R1 to peer with Router R2, and assign the local address and neighbor autonomous system to the ebgp BGP group.

```
[edit protocols]
user@R1# set bgp group ebgp type external
user@R1# set bgp group ebgp neighbor 8.42.1.104 local-address 8.42.1.102
user@R1# set bgp group ebgp neighbor 8.42.1.104 peer-as 65534
```

9. Include the BGP-TE signaling NLRI to the ebgp BGP group.

```
[edit protocols]
user@R1# set bgp group ebgp family traffic-engineering unicast
```

10. Enable passive traffic-engineering on the inter-AS link.


```
[edit protocols]
```

```
user@R1# set isis interface ge-0/0/1.0 passive remote-node-iso 0102.5502.4211
```

```
user@R1# set isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.104
```

11. Enable OSPF on the interface connecting Router R1 to Router R0 and on the loopback interface of Router R1, and enable traffic engineering capabilities.

```
[edit protocols]
```

```
user@R1# set ospf traffic-engineering
```

```
user@R1# set ospf area 0.0.0.0 interface lo0.0
```

```
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
```

12. Enable passive traffic-engineering on the inter-AS link.

```
[edit protocols]
```

```
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-id 8.42.1.104
```

```
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-router-id 10.255.105.139
```

13. Configure policies to accept traffic from BGP-TE NLRI.

```
[edit policy-options]
```

```
user@R1# set policy-statement accept-all from family traffic-engineering
```

```
user@R1# set policy-statement accept-all then accept
```

```
user@R1# set policy-statement nlri2bgp term 1 from family traffic-engineering
```

```
user@R1# set policy-statement nlri2bgp term 1 then accept
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 8.31.1.103/24;
    }
    family iso;
```



```

        family mpls;
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 8.42.1.102/24;
        }
        family iso;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.105.141/32;
        }
    }
}
}

```

```

user@R1# show routing-options
router-id 10.255.105.141;
autonomous-system 65533;

```

```

user@R1# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group ibgp {
        type internal;
        local-address 10.255.105.141;
        family traffic-engineering {
            unicast;

```



```

    }
    export nlri2bgp;
    neighbor 10.255.105.137;
  }
  group ebgp {
    type external;
    family traffic-engineering {
      unicast;
    }
    neighbor 8.42.1.104 {
      local-address 8.42.1.102;
      peer-as 65534;
    }
  }
}
isis {
  interface ge-0/0/1.0 {
    passive {
      remote-node-iso 0102.5502.4211;
      remote-node-id 8.42.1.104;
    }
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface lo0.0;
    interface ge-0/0/0.0;
    interface ge-0/0/1.0 {
      passive {
        traffic-engineering {
          remote-node-id 8.42.1.104;
          remote-node-router-id 10.255.105.139;
        }
      }
    }
  }
}
}

```

```

user@R1# show policy-options
policy-statement accept-all {
  from family traffic-engineering;
  then accept;
}

```



```

policy-statement nlri2bgp {
  term 1 {
    from family traffic-engineering;
    then {
      accept;
    }
  }
}

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure Router R2:

1. Configure the Router R2 interfaces.

```

[edit interfaces]
user@R2# set ge-0/0/0 unit 0 family inet address 8.64.1.104/24
user@R2# set ge-0/0/0 unit 0 family iso
user@R2# set ge-0/0/0 unit 0 family mpls
user@R2# set ge-0/0/1 unit 0 family inet address 8.42.1.104/24
user@R2# set ge-0/0/1 unit 0 family iso
user@R2# set ge-0/0/1 unit 0 family mpls
user@R2# set lo0 unit 0 family inet address 10.255.105.139/32
user@R2# set lo0 unit 0 family iso

```

2. Configure the router ID and autonomous system of Router R2.

```

[edit routing-options]
user@R2# set router-id 10.255.105.139
user@R2# set autonomous-system 65534

```

3. Enable RSVP on all the interfaces of Router R2 (excluding the management interface).

```

[edit routing-options]
user@R2# set rsvp interface all
user@R2# set rsvp interface fxp0.0 disable

```

4. Enable MPLS on all the interfaces of Router R2 (excluding the management interface).


```
[edit routing-options]
user@R2# set mpls interface all
user@R2# set mpls interface fxp0.0 disable
```

5. Enable import of traffic engineering database parameters using the ted2nlri policy.

```
[edit protocols]
user@R2# set mpls traffic-engineering database import policy ted2nlri
```

6. Configure the BGP group for Router R2 to peer with Router R1.

```
[edit protocols]
user@R2# set bgp group ebgp type external
```

7. Include the BGP-TE signaling NLRI to the ebgp BGP group.

```
[edit protocols]
user@R2# set bgp group ebgp family traffic-engineering unicast
```

8. Assign the local address and neighbor autonomous system to the ebgp BGP group.

```
[edit protocols]
user@R2# set bgp group ebgp peer-as 65533
user@R2# set bgp group ebgp neighbor 8.42.1.102
```

9. Enable export of policy nlri2bgp on Router R2.

```
[edit protocols]
user@R2# set bgp group ebgp export nlri2bgp
```

10. Enable IS-IS on the interface connecting Router R2 with Router R3 and the loopback interface of Router R2.

```
[edit protocols]
user@R2# set isis level 1 disable
user@R2# set isis interface ge-0/0/0.0
user@R2# set isis interface lo0.0
```


11. Enable only IS-IS advertising on the interface connecting Router R2 with Router R1.

```
[edit protocols]
user@R2# set isis interface ge-0/0/1.0 passive remote-node-iso 0102.5501.8181
user@R2# set isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.102
```

12. Configure traffic engineering capability on Router R2.

```
[edit protocols]
user@R2# set ospf traffic-engineering
```

13. Enable only OSPF advertisements on the interface connecting Router R2 with Router R1.

```
[edit protocols]
user@R2# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-id 8.42.1.102
user@R2# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-router-id
10.255.105.141
```

14. Configure policies to accept traffic from the BGP-TE NLRI.

```
[edit policy-options]
user@R2# set policy-statement accept-all from family traffic-engineering
user@R2# set policy-statement accept-all then accept
user@R2# set policy-statement nlri2bgp term 1 from family traffic-engineering
user@R2# set policy-statement nlri2bgp term 1 then accept
user@R2# set policy-statement ted2nlri term 1 from protocol isis
user@R2# set policy-statement ted2nlri term 1 from protocol ospf
user@R2# set policy-statement ted2nlri term 1 then accept
user@R2# set policy-statement ted2nlri term 2 then reject
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 8.64.1.104/24;
```



```

    }
    family iso;
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 8.42.1.104/24;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.105.139/32;
    }
    family iso;
  }
}
}

```

```

user@R2# show routing-options
router-id 10.255.105.139;
autonomous-system 65534;

```

```

user@R2# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  traffic-engineering {
    database {
      import {
        policy ted2nlri;
      }
    }
  }
}
interface all;

```



```

interface fxp0.0 {
    disable;
}
}
bgp {
    group ebgp {
        type external;
        family traffic-engineering {
            unicast;
        }
        export nlri2bgp;
        peer-as 65533;
        neighbor 8.42.1.102;
    }
}
isis {
    level 1 disable;
    interface ge-0/0/0.0;
    interface ge-0/0/1.0 {
        passive {
            remote-node-iso 0102.5501.8181;
            remote-node-id 8.42.1.102;
        }
    }
    interface lo0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/0/1.0 {
            passive {
                traffic-engineering {
                    remote-node-id 8.42.1.102;
                    remote-node-router-id 10.255.105.141;
                }
            }
        }
    }
}
}

```

```

user@R2# show policy-options
policy-statement accept-all {
    from family traffic-engineering;
    then accept;
}

```



```

}
policy-statement nlri2bgp {
  term 1 {
    from family traffic-engineering;
    then {
      accept;
    }
  }
}
policy-statement ted2nlri {
  term 1 {
    from protocol [ isis ospf ];
    then accept;
  }
  term 2 {
    then reject;
  }
}

```

Verification

IN THIS SECTION

- [Verifying the BGP Summary Status | 1178](#)
- [Verifying the MPLS LSP Status | 1179](#)
- [Verifying the Isdist.0 Routing Table Entries | 1180](#)
- [Verifying the Traffic Engineering Database Entries | 1184](#)

Verify that the configuration is working properly.

Verifying the BGP Summary Status

Purpose

Verify that BGP is up and running on Routers R0 and R1.

Action

From operational mode, run the **show bgp summary** command.

```
user@R0> show bgp summary
```



```

Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
lsdist.0
              10         10         0         0         0         0
Peer          AS      InPkt   OutPkt   OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.255.105.141 65533      20      14      0      79      5:18
Establ
  lsdist.0: 10/10/10/0

```

From operational mode, run the **show bgp summary** command.

```
user@R1> show bgp summary
```

```

Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
lsdist.0
              10         10         0         0         0         0
Peer          AS      InPkt   OutPkt   OutQ   Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
8.42.1.104     65534      24      17      0      70      6:43
Establ
  lsdist.0: 10/10/10/0
10.255.105.137 65533      15      23      0      79      6:19
Establ
  lsdist.0: 0/0/0/0

```

Meaning

Router R0 is peered with Router R1.

Verifying the MPLS LSP Status

Purpose

Verify the status of the MPLS LSP on Router R0.

Action

From operational mode, run the **show mpls lsp** command.

```
user@R0> show mpls lsp
```



```

Ingress LSP: 1 sessions
To           From           State Rt P      ActivePath      LSPname
10.255.105.135 10.255.105.137 Up      0 *
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The MPLS LSP from Router R0 to Router R3 is established.

Verifying the Isdist.0 Routing Table Entries

Purpose

Verify the Isdist.0 routing table entries on Routers R0, R1, and R2.

Action

From operational mode, run the **show route table Isdist.0** command.

```
user@R0> show route table Isdist.0
```

```

Isdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

NODE { AS:65534 ISO:0102.5502.4211.00 ISIS-L2:0 }/1152
    *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
    AS path: 65534 I, validation-state: unverified
    > to 8.31.1.103 via ge-0/0/0.0
NODE { AS:65534 ISO:0102.5502.4250.00 ISIS-L2:0 }/1152
    *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
    AS path: 65534 I, validation-state: unverified
    > to 8.31.1.103 via ge-0/0/0.0
NODE { AS:65534 ISO:0102.5502.4250.02 ISIS-L2:0 }/1152
    *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
    AS path: 65534 I, validation-state: unverified
    > to 8.31.1.103 via ge-0/0/0.0
NODE { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 OSPF:0 }/1152
    *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
    AS path: 65534 I, validation-state: unverified
    > to 8.31.1.103 via ge-0/0/0.0

```



```

LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.42.1.104 } Remote {
AS:65534 ISO:0102.5501.8181.00 }.{ IPv4:8.42.1.102 } ISIS-L2:0 }/1152

      *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.64.1.104 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
      *[BGP/170] 00:02:03, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4250.00 }.{ IPv4:8.64.1.106 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
      *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4211.00 }.{ } ISIS-L2:0 }/1152
      *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4250.00 }.{ } ISIS-L2:0 }/1152
      *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 }.{ IPv4:8.42.1.104 }
Remote { AS:65534 Area:0.0.0.0 IPv4:10.255.105.141 }.{ IPv4:8.42.1.102 } OSPF:0
}/1152

      *[BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0

```

From operational mode, run the **show route table lsdist.0** command.

user@R1> **show route table lsdist.0**

```

lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

NODE { AS:65534 ISO:0102.5502.4211.00 ISIS-L2:0 }/1152
      *[BGP/170] 00:18:00, localpref 100
      AS path: 65534 I, validation-state: unverified

```



```

> to 8.42.1.104 via ge-0/0/1.0
NODE { AS:65534 ISO:0102.5502.4250.00 ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
NODE { AS:65534 ISO:0102.5502.4250.02 ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
NODE { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 OSPF:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.42.1.104 } Remote {
AS:65534 ISO:0102.5501.8181.00 }.{ IPv4:8.42.1.102 } ISIS-L2:0 }/1152

    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.64.1.104 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:02:19, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4250.00 }.{ IPv4:8.64.1.106 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4211.00 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4250.00 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 }.{ IPv4:8.42.1.104 }
Remote { AS:65534 Area:0.0.0.0 IPv4:10.255.105.141 }.{ IPv4:8.42.1.102 } OSPF:0
}/1152

    *[BGP/170] 00:18:00, localpref 100

```



```
AS path: 65534 I, validation-state: unverified
> to 8.42.1.104 via ge-0/0/1.0
```

From operational mode, run the **show route table lsdist.0** command.

user@R2> **show route table lsdist.0**

```
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

NODE { AS:65534 ISO:0102.5502.4211.00 ISIS-L2:0 }/1152
    *[IS-IS/18] 1d 00:24:39
    Fictitious
NODE { AS:65534 ISO:0102.5502.4250.00 ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
    Fictitious
NODE { AS:65534 ISO:0102.5502.4250.02 ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
    Fictitious
NODE { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 OSPF:0 }/1152
    *[OSPF/10] 1d 00:24:39
    Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.42.1.104 } Remote {
AS:65534 ISO:0102.5501.8181.00 }.{ IPv4:8.42.1.102 } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:58
    Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.64.1.104 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:02:34
    Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4250.00 }.{ IPv4:8.64.1.106 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
    Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4211.00 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
    Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4250.00 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
    Fictitious
```



```

LINK { Local { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 }.{ IPv4:8.42.1.104 }
Remote { AS:65534 Area:0.0.0.0 IPv4:10.255.105.141 }.{ IPv4:8.42.1.102 } OSPF:0
}/1152

*[OSPF/10] 00:20:57
Fictitious

```

Meaning

The routes are appearing in the lsdist.0 routing table.

Verifying the Traffic Engineering Database Entries

Purpose

Verify the traffic engineering database entries on Router R0.

Action

From operational mode, run the **show ted database** command.

user@R0> **show ted database**

```

TED database: 5 ISIS nodes 5 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5501.8168.00(10.255.105.137) Rtr   1046     1      1 OSPF(0.0.0.0)
    To: 8.31.1.101-1, Local: 8.31.1.101, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5501.8181.00                 ---    1033     1      0
0102.5502.4211.00(10.255.105.139) Rtr   3519     2      3 Exported ISIS-L2(1)
    To: 0102.5502.4250.02, Local: 8.64.1.104, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: 0102.5501.8181.00, Local: 8.42.1.104, Remote: 8.42.1.102
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
                                Exported OSPF(2)
    To: 10.255.105.141, Local: 8.42.1.104, Remote: 8.42.1.102
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5502.4250.00(10.255.105.135) Rtr   1033     1      1 Exported ISIS-L2(1)
    To: 0102.5502.4250.02, Local: 8.64.1.106, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5502.4250.02                 Net    1033     2      2 Exported ISIS-L2(1)
    To: 0102.5502.4211.00(10.255.105.139), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```



```

To: 0102.5502.4250.00(10.255.105.135), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                                     Type Age(s) LnkIn LnkOut Protocol
8.31.1.101-1                          Net   1046     2     2 OSPF(0.0.0.0)
  To: 0102.5501.8168.00(10.255.105.137), Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: 10.255.105.141, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                                     Type Age(s) LnkIn LnkOut Protocol
10.255.105.141                        Rtr   1045     2     2 OSPF(0.0.0.0)
  To: 0102.5502.4211.00(10.255.105.139), Local: 8.42.1.102, Remote: 8.42.1.104
    Local interface index: 0, Remote interface index: 0
  To: 8.31.1.101-1, Local: 8.31.1.103, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```

Meaning

The routes are appearing in the traffic engineering database.

Release History Table

Release	Description
17.1R1	Starting with Junos OS Release 17.1R1, link state distribution using BGP is supported on QFX10000 switches.

RELATED DOCUMENTATION

| [Link-State Distribution Using BGP Overview](#) | **1148**

Configuring Link State Distribution Using BGP

You can enable distribution of topology information across multiple areas and autonomous systems (ASs) by extending the BGP protocol to carry link-state information, which was initially acquired using IGP. The IGP protocols have scaling limitations when it comes to distributing large databases. BGP is not only a more scalable vehicle for carrying multi-area and multi-AS topology information, but also provides the policy controls that can be useful for multi-AS topology distribution. The BGP link-state topology information is used for computing paths for MPLS LSPs spanning multiple domains, such as inter-area TE LSP, and providing a scalable and policy-controlled means for external path computing entities, such as ALTO and PCE, to acquire network topology.

Before you begin:

1. Configure the device interfaces.
2. Configure the router ID and autonomous system number for the device.
3. Configure the following protocols:
 - RSVP
 - MPLS
 - IS-IS
 - OSPF

To enable link-state distribution using BGP:

1. Configure an internal BGP group, and assign the local address and neighbor address for the group.

```
[edit protocols]
user@R1# set bgp group internal-group-name type internal
user@R1# set bgp group internal-group-name local-address ip-address
user@R1# set bgp group internal-group-name neighbor ip-address
```

2. Include the BGP-TE signaling network layer reachability information (NLRI) to the internal BGP group.

```
[edit protocols]
user@R1# set bgp group internal-group-name family traffic-engineering unicast
```

3. Enable export of policy on the device.

```
[edit protocols]
user@R1# set bgp group internal-group-name export second-policy-name
```


4. Configure an external BGP group, and assign the local address and neighbor autonomous system to the group.

```
[edit protocols]
user@R1# set bgp group external-group-name type external
user@R1# set bgp group external-group-name neighbor ip-address local-address ip-address
user@R1# set bgp group external-group-name neighbor ip-address peer-as as-number
```

5. Include the BGP-TE signaling NLRI to the external BGP group.

```
[edit protocols]
user@R1# set bgp group external-group-name family traffic-engineering unicast
```

6. In configuration mode, go to the following hierarchy level:

```
[edit]
user@R1# edit policy-options
```

7. Configure policies to accept traffic from the BGP-TE NLRI.

```
[edit policy-options]
user@R1# set policy-statement policy-name from family traffic-engineering
user@R1# set policy-statement policy-name then accept
user@R1# set policy-statement bgp-import-policy term 1 from family traffic-engineering
user@R1# set policy-statement bgp-import-policy term 1 then next-hop self
user@R1# set policy-statement bgp-import-policy term 1 then accept
```

8. On the remote connecting device, configure policy to accept the OSPF and IS-IS traffic.

```
[edit policy-options]
user@R2# set policy-statement bgp-export-policy term 1 from protocol isis
user@R2# set policy-statement bgp-export-policy term 1 from protocol ospf
user@R2# set policy-statement bgp-export-policy term 1 then accept
user@R2# set policy-statement bgp-export-policy term 2 then reject
```

9. Verify and commit the configuration.

For example:

R1

[edit protocols]

```

user@R1# set rsvp interface all
user@R1# set rsvp interface fxp0.0 disable
user@R1# set mpls interface all
user@R1# set mpls interface fxp0.0 disable
user@R1# set bgp group ibgp type internal
user@R1# set bgp group ibgp local-address 10.255.105.141
user@R1# set bgp group ibgp family traffic-engineering unicast
user@R1# set bgp group ibgp export nlri2bgp
user@R1# set bgp group ibgp neighbor 10.255.105.137
user@R1# set bgp group ebgp type external
user@R1# set bgp group ebgp family traffic-engineering unicast
user@R1# set bgp group ebgp neighbor 8.42.1.104 local-address 8.42.1.102
user@R1# set bgp group ebgp neighbor 8.42.1.104 peer-as 65534
user@R1# set isis interface ge-0/0/1.0 passive remote-node-iso 0102.5502.4211
user@R1# set isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.104
user@R1# set ospf traffic-engineering
user@R1# set ospf area 0.0.0.0 interface lo0.0
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-id 8.42.1.104
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering remote-node-router-id
10.255.105.139

```

[edit policy-options]

```

user@R1# set policy-statement accept-all from family traffic-engineering
user@R1# set policy-statement accept-all then accept
user@R1# set policy-statement nlri2bgp term 1 from family traffic-engineering
user@R1# set policy-statement nlri2bgp term 1 then next-hop self
user@R1# set policy-statement nlri2bgp term 1 then accept

```

[edit]

```

user@R1# commit
commit complete

```

R2

[edit policy-options]

```

user@R2# set policy-statement accept-all from family traffic-engineering
user@R2# set policy-statement accept-all then accept
user@R2# set policy-statement nlri2bgp term 1 from family traffic-engineering

```



```

user@R2# set policy-statement nlri2bgp term 1 then next-hop self
user@R2# set policy-statement nlri2bgp term 1 then accept
user@R2# set policy-statement ted2nlri term 1 from protocol isis
user@R2# set policy-statement ted2nlri term 1 from protocol ospf
user@R2# set policy-statement ted2nlri term 1 then accept
user@R2# set policy-statement ted2nlri term 2 then reject

```

```

[edit]
user@R2# commit
commit complete

```

RELATED DOCUMENTATION

[Link-State Distribution Using BGP Overview | 1148](#)

[Example: Configuring Link State Distribution Using BGP | 1163](#)

Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages

IN THIS SECTION

- [PathErr Messages | 1190](#)
- [Identifying the Problem Link | 1191](#)
- [Configuring the Router to Improve Traffic Engineering Database Accuracy | 1191](#)

An essential element of RSVP-based traffic engineering is the traffic engineering database. The traffic engineering database contains a complete list of all network nodes and links participating in traffic engineering, and a set of attributes each of those links can hold. (For more information about the traffic engineering database, see [“Constrained-Path LSP Computation” on page 485](#).) One of the most important link attributes is bandwidth.

Bandwidth availability on links changes quickly as RSVP LSPs are established and terminated. It is likely that the traffic engineering database will develop inconsistencies relative to the real network. These inconsistencies cannot be fixed by increasing the rate of IGP updates.

Link availability can share the same inconsistency problem. A link that becomes unavailable can break all existing RSVP LSPs. However, its unavailability might not readily be known by the network.

When you configure the **rsvp-error-hold-time** statement, a source node (ingress of an RSVP LSP) learns from the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes. Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the traffic engineering database and the network.

You can control the frequency of IGP updates by using the **update-threshold** statement. See [“Configuring the RSVP Update Threshold on an Interface” on page 833](#).

This section discusses the following topics:

PathErr Messages

PathErr messages report a wide variety of problems by means of different code and subcode numbers. You can find a complete list of these PathErr messages in RFC 2205, *Resource Reservation Protocol (RSVP), Version 1, Functional Specification* and RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*.

When you configure the **rsvp-error-hold-time** statement, two categories of PathErr messages, which specifically represent link failures, are examined:

- Link bandwidth is low for this LSP: Requested bandwidth unavailable—code 1, subcode 2

This type of PathErr message represents a global problem that affects all LSPs transiting the link. They indicate that the actual link bandwidth is lower than that required by the LSP, and that it is likely that the bandwidth information in the traffic engineering database is an overestimate.

When this type of error is received, the available link bandwidth is reduced in the local traffic engineering database, affecting all future LSP computations.

- Link unavailable for this LSP:
 - Admission Control failure—code 1, any subcode except 2
 - Policy Control failures—code 2
 - Service Preempted—code 12
 - Routing problem—no route available toward destination—code 24, subcode 5

These types of PathErr messages are generally pertinent to the specified LSP. The failure of this LSP does not necessarily imply that other LSPs could also fail. These errors can indicate maximum transfer unit (MTU) problems, service preemption (either manually initiated by the operator or by another LSP with a higher priority), that a next-hop link is down, that a next-hop neighbor is down, or service rejection because of policy considerations. It is best to route this particular LSP away from the link.

Identifying the Problem Link

Each PathErr message includes the sender's IP address. This information is propagated unchanged toward the ingress router. A lookup in the traffic engineering database can identify the node that originated the PathErr message.

Each PathErr message carries enough information to identify the RSVP session that triggered the message. If this is a transit router, it simply forwards the message. If this router is the ingress router (for this RSVP session), it has the complete list of all nodes and links the session should traverse. Coupled with the originating node information, the link can be uniquely identified.

Configuring the Router to Improve Traffic Engineering Database Accuracy

To improve the accuracy of the traffic engineering database, configure the **rsvp-error-hold-time** statement. When this statement is configured, a source node (ingress of an RSVP LSP) learns from the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes. Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages also are used to update traffic engineering database bandwidth information, reducing inconsistencies between the traffic engineering database and the network.

To configure how long MPLS should remember RSVP PathErr messages and consider them in CSPF computation, include the **rsvp-error-hold-time** statement:

```
rsvp-error-hold-time seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The time can be a value from 1 to 240 seconds. The default is 25 seconds. Configuring a value of 0 disables the monitoring of PathErr messages.

Configuring DiffServ-Aware Traffic Engineering to Achieve Service Level Guarantees on an MPLS network

IN THIS CHAPTER

- DiffServ-Aware Traffic Engineering Introduction | 1194
- DiffServ-Aware Traffic Engineering Standards | 1194
- DiffServ-Aware Traffic Engineering Terminology | 1194
- DiffServ-Aware Traffic Engineering Features | 1195
- DiffServ-Aware Traffic Engineered LSPs | 1196
- DiffServ-Aware Traffic Engineered LSPs Overview | 1196
- DiffServ-Aware Traffic Engineered LSPs Operation | 1197
- Multiclass LSP Overview | 1198
- Multiclass LSPs | 1198
- Establishing a Multiclass LSP on the Differentiated Services Domain | 1199
- Configuring Routers for DiffServ-Aware Traffic Engineering | 1199
- LSP Bandwidth Oversubscription Overview | 1204
- LSP Size Oversubscription | 1205
- LSP Link Size Oversubscription | 1206
- Class Type Oversubscription and Local Oversubscription Multipliers | 1206
- Class Type Bandwidth and the LOM | 1207
- LOM Calculation for the MAM and Extended MAM Bandwidth Models | 1207
- LOM Calculation for the Russian Dolls Bandwidth Model | 1208
- Example: LOM Calculation | 1208
- Configuring the Bandwidth Subscription Percentage for LSPs | 1210
- Configuring LSPs for DiffServ-Aware Traffic Engineering | 1212
- Configuring Multiclass LSPs | 1215

DiffServ-Aware Traffic Engineering Introduction

Differentiated Services (DiffServ)-aware traffic engineering provides a way to guarantee a specified level of service over an MPLS network. The routers providing DiffServ-aware traffic engineering are part of a differentiated services network domain. All routers participating in a differentiated services domain must have DiffServ-aware traffic engineering enabled.

To help ensure that the specified service level is provided, it is necessary to ensure that no more than the amount of traffic specified is sent over the differentiated services domain. You can accomplish this goal by configuring a policer to police or rate-limit the volume of traffic transiting the differentiated service domain. For more information about how to configure policers for label-switched paths (LSPs), see [“Configuring Policers for LSPs” on page 121](#).

This feature can help to improve the quality of Internet services such as voice over IP (VoIP). It also makes it possible to better emulate an Asynchronous Transfer Mode (ATM) circuit over an MPLS network.

DiffServ-Aware Traffic Engineering Standards

The following RFCs provide information on DiffServ-aware traffic engineering and multiclass LSPs:

- RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*
- RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 4124, *Protocol Extensions for Support of Differentiated-Service-Aware MPLS Traffic Engineering*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS*

These RFCs are available on the IETF website at <http://www.ietf.org/>.

DiffServ-Aware Traffic Engineering Terminology

B

Bandwidth model	The bandwidth model determines the values of the available bandwidth advertised by the interior gateway protocols (IGPs).
------------------------	---

C

CAC	Call admission control (CAC) checks to ensure there is adequate bandwidth on the path before the LSP is established. If the bandwidth is insufficient, the LSP is not established and an error is reported.
------------	---

Class type	A collection of traffic flows that is treated equivalently in a differentiated services domain. A class type maps to a queue and is much like a class-of-service (CoS) forwarding class in concept. It is also known as a traffic class.
D	
Differentiated Services	Differentiated Services make it possible to give different treatment to traffic based on the EXP bits in the MPLS header. Traffic must be marked appropriately and CoS must be configured.
Differentiated Services domain	The routers in a network that have Differentiated Services enabled.
DiffServ-aware traffic engineering	A type of constraint-based routing. It can enforce different bandwidth constraints for different classes of traffic. It can also do CAC on each traffic engineering class when an LSP is established.
M	
MAM	The maximum allocation bandwidth constraint model divides the available bandwidth between the different classes. Sharing of bandwidth between the class types is not allowed.
Multiclass LSP	A multiclass LSP functions like a standard LSP, but it also allows you to reserve bandwidth from multiple class types. The EXP bits of the MPLS header are used to distinguish between class types.
R	
RDM	The Russian dolls bandwidth constraint model makes efficient use of bandwidth by allowing the class types to share bandwidth.
T	
Traffic engineering class	A paired class type and priority.
Traffic engineering class map	A map between the class types, priorities, and traffic engineering classes. The traffic engineering class mapping must be consistent across the Differentiated Services domain.

DiffServ-Aware Traffic Engineering Features

DiffServ-aware traffic engineering provides the following features:

- Traffic engineering at a per-class level rather than at an aggregate level
- Different bandwidth constraints for different class types (traffic classes)
- Different queuing behaviors per class, allowing the router to forward traffic based on the class type

In comparison, standard traffic engineering does not consider CoS, and it completes its work on an aggregate basis across all Differentiated Service classes.

DiffServ-aware traffic engineering provides the following advantages:

- Traffic engineering can be performed on a specific class type instead of at the aggregate level.
- Bandwidth constraints can be enforced on each specific class type.
- It forwards traffic based on the EXP bits.

This makes it possible to guarantee service and bandwidth across an MPLS network. With DiffServ-aware traffic engineering, among other services, you can provide ATM circuit emulation, VoIP, and a guaranteed bandwidth service.

The following describes how the IGP, Constrained Shortest Path First (CSPF), and RSVP participate in DiffServ-aware traffic engineering:

- The IGP can advertise the unreserved bandwidth for each traffic engineering class to the other members of the differentiated services domain. The traffic engineering database stores this information.
- A CSPF calculation is performed considering the bandwidth constraints for each class type. If all the constraints are met, the CSPF calculation is considered successful.
- When RSVP signals an LSP, it requests bandwidth for specified class types.

DiffServ-Aware Traffic Engineered LSPs

A DiffServ-aware traffic engineered LSP is an LSP configured to reserve bandwidth for one of the supported class types and to carry traffic for that class type. The following sections discuss this type of LSPs:

- [DiffServ-Aware Traffic Engineered LSPs Overview on page 1196](#)
- [DiffServ-Aware Traffic Engineered LSPs Operation on page 1197](#)

DiffServ-Aware Traffic Engineered LSPs Overview

A DiffServ-aware traffic engineered LSP is an LSP configured with a bandwidth reservation for a specific class type. This LSP can carry traffic for a single class type. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

The class type must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network. You can unambiguously map a class type to a queue. On each node router, the class-of-service queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

For more information about topics related to LSPs and DiffServ-aware traffic engineering, see the following:

- For forwarding classes and class of service, see the *Class of Service User Guide (Routers and EX9200 Switches)*.
- For EXP bits, see [“MPLS Label Allocation” on page 420](#).
- For differentiated services, see RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*.
- For information about how the IGPs and RSVP have been modified to support Differentiated Services-aware MPLS traffic engineering, see RFC 4124, *Protocol Extensions for Support of Differentiated-Service-Aware MPLS Traffic Engineering*.

DiffServ-Aware Traffic Engineered LSPs Operation

When configuring a DiffServ-aware traffic engineered LSP, you specify the class type and the bandwidth associated with it. The following occurs when an LSP is established with bandwidth reservation from a specific class type:

1. The IGPs advertise how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for an LSP, CSPF is used to ensure that the bandwidth constraints are met for the class type carried by the LSP at the specified priority level.

CSPF also checks to ensure that the bandwidth model is configured consistently on each router participating in the LSP. If the bandwidth model is inconsistent, CSPF does not compute the path (except for LSPs from class type ct0).

3. Once a path is found, RSVP signals the LSP using the Classtype object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up.

An LSP that requires bandwidth from a particular class (except class type ct0) cannot be established through routers that do not understand the Classtype object. Preventing the use of routers that do not understand the Classtype object helps to ensure consistency throughout the Differentiated Services domain by preventing the LSP from using a router that cannot support Differentiated Services.

By default, LSPs are signaled with setup priority 7 and holding priority 0. An LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both LSPs configured for DiffServ-aware traffic engineering and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings (either by remarking the EXP settings or by assuming that the traffic arrived with the correct EXP settings from the upstream router).

Multiclass LSP Overview

A multiclass LSP is an LSP that can carry several class types. One multiclass LSP can be used to support up to four class types. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

Once a multiclass LSP is configured, traffic from all of the class types can:

- Follow the same path
- Be rerouted along the same path
- Be taken down at the same time

Class types must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network.

You can unambiguously map a class type to a queue. On each node router, the CoS queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

The combination of a class type and a priority level forms a traffic engineering class. The IGPs can advertise up to eight traffic engineering classes for each link.

For more information about the EXP bits, see [“MPLS Label Allocation” on page 420](#).

For more information about forwarding classes, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

Multiclass LSPs

Multiclass LSPs function like standard LSPs, but they also allow you to configure multiple class types with guaranteed bandwidth. The EXP bits of the MPLS header are used to distinguish between class types. Multiclass LSPs can be configured for a variety of purposes. For example, you can configure a multiclass LSP to emulate the behavior of an ATM circuit. An ATM circuit can provide service-level guarantees to a class type. A multiclass LSP can provide a similar guaranteed level of service.

The following sections discuss multiclass LSPs:

- [Multiclass LSP Overview on page 1198](#)
- [Establishing a Multiclass LSP on the Differentiated Services Domain on page 1199](#)

Establishing a Multiclass LSP on the Differentiated Services Domain

The following occurs when a multiclass LSP is established on the differentiated services domain:

1. The IGP advertises how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for a multiclass LSP, CSPF is used to ensure that the constraints are met for all the class types carried by the multiclass LSP (a set of constraints instead of a single constraint).
3. Once a path is found, RSVP signals the LSP using an RSVP object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up. The RSVP object is a hop-by-hop object. Multiclass LSPs cannot be established through routers that do not understand this object. Preventing routers that do not understand the RSVP object from carrying traffic helps to ensure consistency throughout the differentiated services domain by preventing the multiclass LSP from using a router that is incapable of supporting differentiated services.

By default, multiclass LSPs are signaled with setup priority 7 and holding priority 0. A multiclass LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both multiclass LSPs and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings.

Configuring Routers for DiffServ-Aware Traffic Engineering

IN THIS SECTION

- [Configuring the Bandwidth Model | 1201](#)
- [Configuring Traffic Engineering Classes | 1202](#)
- [Configuring Class of Service for DiffServ-Aware Traffic Engineering | 1204](#)

To configure DiffServ-aware traffic engineering, include the **diffserv-te** statement:

```
diffserv-te {
  bandwidth-model {
    extended-mam;
    mam;
    rdm;
```



```

    }
    te-class-matrix {
        traffic-class {
            tnumber {
                priority priority;
                traffic-class ctnumber priority priority;
            }
        }
    }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

You must include the **diffserv-te** statement in the configuration on all routers participating in the Differentiated Services domain. However, you are not required to configure the traffic engineering class matrix (by including the **te-class-matrix** statement at the [edit protocols mpls **diffserv-te**] or [edit logical-systems *logical-system-name* protocols mpls **diffserv-te**] hierarchy level).

NOTE: To prevent the possibility of an incorrect configuration when migrating to Diffserv-aware traffic engineering, a policy control failure error might be triggered if there is conflict between the old LSPs and the newly configured TE-class matrix.

An old node might request an LSP with setup and hold priorities in such a way that the combination of the ct0 class and the priority does not match with the configured TE-class matrix. All LSPs on the router that are configured prior to configuring diffserv-aware traffic engineering are designated as being from class ct0.

The error appears in the RSVP tracing logs as a **Session preempted** error. For the router where the error originates, the error could appear as follows:

```
Jun 17 16:35:59 RSVP error for session 10.255.245.6(port/tunnel ID 31133)
  Proto 0: (class ct0, priority 2) is not a valid TE-class Jun 17 16:35:59
  RSVP originate PathErr 192.168.37.22->192.168.37.23 Session preempted
```

For the router receiving the error, the error can appear as follows:

```
Jun 17 16:37:51 RSVP recv PathErr 192.168.37.22->192.168.37.23 Session
preempted LSP to-f(2/31133)
```

To configure DiffServ-aware traffic engineering, complete the procedures in the following sections:

Configuring the Bandwidth Model

You must configure a bandwidth model on all routers participating in the Differentiated Services domain. The bandwidth models available are MAM, extended MAM, and RDM:

- Maximum allocation bandwidth constraints model (MAM)—Defined in RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.
- Extended MAM—A proprietary bandwidth model that behaves much like standard MAM. If you configure multiclass LSPs, you must configure the extended MAM bandwidth model.
- Russian-dolls bandwidth allocation model (RDM)—Makes efficient use of bandwidth by allowing the class types to share bandwidth. RDM is defined in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.

To configure a bandwidth model, include the **bandwidth-model** statement and specify one of the bandwidth model options:


```
bandwidth-model {
  extended-mam;
  mam;
  rdm;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [diffserv-te](#)]
- [edit logical-systems *logical-system-name* protocols mpls [diffserv-te](#)]

NOTE: If you change the bandwidth model on an ingress router, all the LSPs enabled on the router are taken down and resignaled.

Configuring Traffic Engineering Classes

Configuring traffic engineering classes is optional. [Table 27 on page 1202](#) shows the default values for everything in the traffic engineering class matrix. The default mapping is expressed in terms of the default forwarding classes defined in the CoS configuration.

Table 27: Default Values for the Traffic Engineering Class Matrix

Traffic Engineering Class	Class Type	Queue	Priority
te0	ct0	0	7
te1	ct1	1	7
te2	ct2	2	7
te3	ct3	3	7
te4	ct0	0	0
te5	ct1	1	0
te6	ct2	2	0
te7	ct3	3	0

If you want to override the default mappings, you can configure traffic engineering classes 0 through 7. For each traffic engineering class, you configure a class type (or queue) from 0 through 3. For each class type, you configure a priority from 0 through 7.

To configure traffic engineering classes explicitly, include the **te-class-matrix** statement:

```
te-class-matrix {
  tnumber {
    priority priority;
    traffic-class {
      ctnumber priority priority;
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **diffserv-te**]
- [edit logical-systems *logical-system-name* protocols mpls **diffserv-te**]

The following example shows how to configure traffic engineering class **te0** with class type **ct1** and a priority of 4:

```
[edit protocols mpls diffserv-te]
te-class-matrix {
  te0 traffic-class ct1 priority 4;
}
```

NOTE: If you explicitly configure a value for one of the traffic engineering classes, all the default values in the traffic engineering class matrix are dropped.

When you explicitly configure traffic engineering classes, you must also configure a bandwidth model; otherwise, the configuration commit operation fails.

Requirements and Limitations for the Traffic Engineering Class Matrix

When you configure a traffic engineering class matrix, be aware of the following requirements and limitations:

- A mapping configuration is local and affects only the router on which it is configured. It does not affect other systems participating in the differentiated services domain. However, for a Differentiated Services domain to function properly, you need to configure the same traffic engineering class matrix on all the routers participating in the same domain.

- When explicitly configuring traffic engineering classes, you must configure the classes in sequence (**te0**, **te1**, **te2**, **te3**, and so on); otherwise, the configuration commit operation fails.

The first traffic engineering class you configure must be **te0**; otherwise, the configuration commit operation fails.

Configuring Class of Service for DiffServ-Aware Traffic Engineering

To configure DiffServ-aware traffic engineering, you must also configure class of service. The following example illustrates a class-of-service configuration that would allocate 25 percent of the link bandwidth to each class:

```
class-of-service {
  interfaces {
    all {
      scheduler-map simple-map;
    }
  }
  scheduler-maps {
    simple-map {
      forwarding-class assured-forwarding scheduler simple_sched;
      forwarding-class best-effort scheduler simple_sched;
      forwarding-class network-control scheduler simple_sched;
      forwarding-class expedited-forwarding scheduler simple_sched;
    }
  }
  schedulers {
    simple_sched {
      transmit-rate percent 25;
      buffer-size percent 25;
    }
  }
}
```

For more information on how to configure class of service, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

LSP Bandwidth Oversubscription Overview

LSPs are established with bandwidth reservations configured for the maximum amount of traffic you expect to traverse the LSP. Not all LSPs carry the maximum amount of traffic over their links at all times. For example, even if the bandwidth for link A has been completely reserved, actual bandwidth might still

be available but not currently in use. This excess bandwidth can be used by allowing other LSPs to also use link A, oversubscribing the link. You can oversubscribe the bandwidth configured for individual class types or specify a single value for all of the class types using an interface.

You can use oversubscription to take advantage of the statistical nature of traffic patterns and to permit higher utilization of links.

The following examples describe how you might use bandwidth oversubscription and undersubscription:

- Use oversubscription on class types where peak periods of traffic do not coincide in time.
- Use oversubscription of class types carrying best-effort traffic. You take the risk of temporarily delaying or dropping traffic in exchange for making better utilization of network resources.
- Give different degrees of oversubscription or undersubscription of traffic for the different class types. For instance, you configure the subscription for classes of traffic as follows:
 - Best effort—**ct0 1000**
 - Voice—**ct3 1**

When you undersubscribe a class type for a multiclass LSP, the total demand of all RSVP sessions is always less than the actual capacity of the class type. You can use undersubscription to limit the utilization of a class type.

The bandwidth oversubscription calculation occurs on the local router only. Because no signaling or other interaction is required from other routers in the network, the feature can be enabled on individual routers without being enabled or available on other routers which might not support this feature. Neighboring routers do not need to know about the oversubscription calculation, they rely on the IGP.

The following sections describe the types of bandwidth oversubscription available in the Junos OS:

- [LSP Size Oversubscription on page 1205](#)
- [LSP Link Size Oversubscription on page 1206](#)
- [Class Type Oversubscription and Local Oversubscription Multipliers on page 1206](#)

LSP Size Oversubscription

For LSP size oversubscription, you simply configure less bandwidth than the peak rate expected for the LSP. You also might need to adjust the configuration for automatic policers. Automatic policers manage the traffic assigned to an LSP, ensuring that it does not exceed the configured bandwidth values. LSP size oversubscription requires that the LSP can exceed its configured bandwidth allocation.

Policing is still possible. However, the policer must be manually configured to account for the maximum bandwidth planned for the LSP, rather than for the configured value.

LSP Link Size Oversubscription

You can increase the maximum reservable bandwidth on the link and use the inflated values for bandwidth accounting. Use the **subscription** statement to oversubscribe the link. The configured value is applied to all class type bandwidth allocations on the link. For more information about link size oversubscription, see [“Configuring the Bandwidth Subscription Percentage for LSPs” on page 1210](#).

Class Type Oversubscription and Local Oversubscription Multipliers

Local oversubscription multipliers (LOMs) allow different oversubscription values for different class types. LOMs are useful for networks where the oversubscription ratio needs to be configured differently on different links and where oversubscription values are required for different classes. You might use this feature to oversubscribe class types handling best-effort traffic, but use no oversubscription for class types handling voice traffic. An LOM is calculated locally on the router. No information related to an LOM is signaled to other routers in the network.

An LOM is configurable on each link and for each class type. The per-class type LOM allows you to increase or decrease the oversubscription ratio. The per-class-type LOM is factored into all local bandwidth accounting for admission control and IGP advertisement of unreserved bandwidths.

The LOM calculation is tied to the bandwidth model (MAM, extended MAM, and Russian dolls) used, because the effect of oversubscription across class types must be accounted for accurately.

NOTE: All LOM calculations are performed by the Junos OS and require no user intervention.

The formulas related to the oversubscription of class types are described in the following sections:

- [Class Type Bandwidth and the LOM on page 1207](#)
- [LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 1207](#)
- [LOM Calculation for the Russian Dolls Bandwidth Model on page 1208](#)
- [Example: LOM Calculation on page 1208](#)

Class Type Bandwidth and the LOM

The following formula expresses the relationship between the bandwidth of the class type and the LOM. The normalized bandwidth of the class type (N_B) is equal to the reserved bandwidth of the class type (R_B) divided by the LOM of the class type (L_C):

$$N_B = R_B / L_C$$

When calculating available bandwidth, you need to subtract the normalized bandwidth from the relevant bandwidth constraint.

NOTE: When using an LOM, values advertised for the available bandwidth might be larger than the bandwidth constraint values. However, the values advertised in the maximum link bandwidth advertisement are not affected by local oversubscription.

LOM Calculation for the MAM and Extended MAM Bandwidth Models

The following formulas show how the LOM is calculated for the MAM and extended MAM bandwidth models.

$$\text{Unreserved TE-Class}(i) = \text{LOM}_c \times [\text{BC}_c - \text{SUM} (\text{Normalized} (\text{CT}_c, q))] \text{ for } q \leq p$$

Or

$$\text{Unreserved TE-Class}(i) = (\text{LOM}_c \times \text{BC}_c) - \text{SUM} (\text{Reserved} (\text{CT}_c, q)) \text{ for } q \leq p$$

where:

- LOM_c —LOM for class type c .
- BC_c —Bandwidth constraint for class type c .
- CT_c —Class type c .
- $\text{TE-Class}(i) \leftrightarrow (\text{CT}_c, \text{preemption } p)$ in the configured TE-Class mapping.

LOM Calculation for the Russian Dolls Bandwidth Model

The following formulas show how the LOM is calculated for the Russian dolls bandwidth model:

```
Unreserved TE-Class (i) = LOMc x MIN [
    [ BCc - SUM ( Normalized (CTb, q) ) ] for q ≤ p and c ≤ b ≤ 7,
    . . .
    [ BC0 - SUM ( Normalized (CTb, q) ) ] for q ≤ p and 0 ≤ b ≤ 7,
]
```

where:

- **LOM_c**—LOM for class type **c**.
- **BC_c**—Bandwidth constraint for class type **c**.
- **TE-Class(i)** <- ->(CT_c , preemption **p**) in the configured TE-Class mapping.

Note that the impact of an LSP on the unreserved bandwidth of a class type does not depend only on the LOM for that class type—it also depends on the LOM for the class type of the LSP.

Example: LOM Calculation

The following example illustrates how an LOM calculation is made for four classes of traffic: **ct0**, **ct1**, **ct2**, and **ct3**.

The class types have been assigned the following values:

```
ct0 = 40
ct1 = 30
ct2 = 20
ct3 = 10
```

These class type values yield the following bandwidth constraints:

```
BC0 = (ct3 + ct2 + ct1 + ct0) = 100
BC1 = (ct3 + ct2 + ct1) = 60
BC2 = (ct3 + ct2) = 30
BC3 = (ct3) = 10
```


LSPs from class type **ct0** can take up to 100 percent of bandwidth on the link. LSPs from class type **ct1** can take up to 60 percent of the bandwidth on the link, and so on.

If you assume for this example that the class types have the following LOM values:

```
LOM(ct0) = 8
LOM(ct1) = 4
LOM(ct2) = 2
LOM(ct3) = 1
```

In the absence of any other reservation, LSPs from class type **ct0** can take up to 800 percent of the available bandwidth ($8 \times 100 = 800$). In the absence of any other reservation, LSPs from class type **ct1** can take up to 240 percent of the available bandwidth ($4 \times 60 = 240$), and so on.

The maximum amount of bandwidth that can be reserved is:

```
ct0 = LOM(ct0) x BC0 = 800
ct1 = LOM(ct1) x BC1 = 240
ct2 = LOM(ct2) x BC2 = 60
ct3 = LOM(ct3) x BC3 = 10
```

For the undersubscribed class type **ct3**, the maximum reservable bandwidth is the same as the bandwidth constraint. For the overbooked class types, these values are not the values of the bandwidth constraint-taking into account the oversubscription for each class type separately. The oversubscription per class type in the sum is not taken into account because ultimately the entire bandwidth constraint can be filled with the bandwidth reservation of just one class type, so you have to account for that class type's bandwidth oversubscription only.

When calculating the available bandwidth for **CTc**, you need to express reservations from other classes as if they were from **CTc**. The reservation from class **ctx** is normalized with the LOM of **ctx**, but it is then multiplied by the LOM of **CTc**.

For the previous example, assume that **LSP1** has class type **ct3** configured with bandwidth of **10** and a priority of **0**.

The values for the reservable bandwidth will be:

```
ct0 = 8 x (100 - 10) = 720
ct1 = 4 x min((100-10), (60-10)) = 200
ct2 = 2 x min((100-10), (60-10), (30-10)) = 40
ct3 = 1 x min((100-10), (60-10), (30-10), (10-10)) = 0
```

These numbers can be rationalized as follows: the normalized reservation is 10 percent. If this bandwidth came from class type **ct0**, it would be equivalent to an overbooked reservation of 80 percent. You can see that 720 percent ($800 - 80 = 720$) of the bandwidth remains available for other LSPs.

Configuring the Bandwidth Subscription Percentage for LSPs

By default, RSVP allows all of a class type's bandwidth (100 percent) to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

If you want to oversubscribe or undersubscribe all of the class types on an interface using the same percentage bandwidth, configure the percentage using the **subscription** statement:

```
subscription percentage;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

To undersubscribe or oversubscribe the bandwidth for each class type, configure a percentage for each class type (**ct0**, **ct1**, **ct2**, and **ct3**) option for the **subscription** statement. When you oversubscribe a class type, an LOM is applied to calculate the actual bandwidth reserved. See [“Class Type Oversubscription and Local Oversubscription Multipliers” on page 1206](#) for more information.

```
subscription {  
    ct0 percentage;  
    ct1 percentage;  
    ct2 percentage;  
    ct3 percentage;  
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

percentage is the percentage of class type bandwidth that RSVP allows to be used for reservations. It can be a value from 0 through 65,000 percent. If you specify a value greater than 100, you are oversubscribing the interface or class type.

The value you configure when you oversubscribe a class type is a percentage of the class type bandwidth that can actually be used. The default subscription value is 100 percent.

You can use the **subscription** statement to disable new RSVP sessions for one or more class types. If you configure a percentage of 0, no new sessions (including those with zero bandwidth requirements) are permitted for the class type.

Existing RSVP sessions are not affected by changing the subscription factor. To clear an existing session, issue the **clear rsvp session** command. For more information on the **clear rsvp session** command, see the [CLI Explorer](#).

Constraints on Configuring Bandwidth Subscription

Be aware of the following issues when configuring bandwidth subscription:

- If you configure bandwidth constraints at the **[edit class-of-service interface *interface-name*]** hierarchy level, they override any bandwidth configuration you specify at the **[edit protocols rsvp interface *interface-name* bandwidth]** hierarchy level for Diffserv-TE. Also note that either of the CoS or RSVP bandwidth constraints can override the interface hardware bandwidth constraints.
- If you configure a bandwidth subscription value for a specific interface that differs from the value configured for all interfaces (by including different values for the **subscription** statement at the **[edit protocols rsvp interface *interface-name*]** and **[edit protocols rsvp interface all]** hierarchy levels), the interface-specific value is used for that interface.
- You can configure subscription for each class type only if you also configure a bandwidth model. If no bandwidth model is configured, the commit operation fails with the following error message:

```
user@host# commit check
```

```
[edit protocols rsvp interface all]
  'subscription'
  RSVP: Must have a diffserv-te bandwidth model configured when configuring
  subscription per traffic class.
  error: configuration check-out failed
```

- You cannot include the **subscription** statement both in the configuration for a specific class type and the configuration for the entire interface. The commit operation fails with the following error message:

```
user@host# commit check
```

```
[edit protocols rsvp interface all]
  'subscription'
    RSVP: Cannot configure both link subscription and per traffic class
  subscription.
  error: configuration check-out failed
```


Configuring LSPs for DiffServ-Aware Traffic Engineering

IN THIS SECTION

- [Configuring Class of Service for the Interfaces | 1213](#)
- [Configuring IGP | 1213](#)
- [Configuring Traffic-Engineered LSPs | 1213](#)
- [Configuring Policing for LSPs | 1214](#)
- [Configuring Fast Reroute for Traffic-Engineered LSPs | 1214](#)

You must configure the Differentiated Services domain (see [“Configuring Routers for DiffServ-Aware Traffic Engineering” on page 1199](#)) before you can enable DiffServ-aware traffic engineering for LSPs. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in the LSP configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the LSP to function properly.

NOTE: You must configure either MAM or RDM as the bandwidth model when you configure DiffServ-aware traffic engineering for LSPs. See [“Configuring the Bandwidth Model” on page 1201](#).

The actual data transmitted over this Differentiated Services domain is carried by an LSP. Each LSP relies on the EXP bits of the MPLS packets to enable DiffServ-aware traffic engineering. Each LSP can carry traffic for a single class type.

All the routers participating in the LSP must be Juniper Networks routers running Junos OS Release 6.3 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the Junos OS. However, the DiffServ-aware traffic engineering LSP cannot traverse these routers.

NOTE: You cannot simultaneously configure multiclass LSPs and DiffServ-aware traffic engineering LSPs on the same router.

To enable DiffServ-aware traffic engineering for LSPs, you need to configure the following:

Configuring Class of Service for the Interfaces

The existing class-of-service (CoS) infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to accomplish this are configured using the existing Junos OS CoS features.

NOTE: The Junos OS does not support CoS on ATM interfaces.

For information about how to configure CoS, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

Configuring IGP

You can configure either IS-IS or OSPF as the IGP. The IS-IS and OSPF configurations for routers supporting LSPs are standard. For information about how to configure these protocols, see the *Junos OS Routing Protocols Library*.

Configuring Traffic-Engineered LSPs

You configure an LSP by using the standard LSP configuration statements and procedures. To configure DiffServ-aware traffic engineering for the LSP, specify a class type bandwidth constraint by including the **bandwidth** statement:

```
label-switched-path lsp-name {
  bandwidth {
    ctnumber bps;
  }
}
```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for this statement.

If you do not specify a bandwidth for a class type, **ct0** is automatically specified as the queue for the LSP. You can configure only one class type for each LSP, unlike multiclass LSPs.

The class type statements specify bandwidth (in bits per second) for the following classes:

- **ct0**—Bandwidth reserved for class 0
- **ct1**—Bandwidth reserved for class 1
- **ct2**—Bandwidth reserved for class 2
- **ct3**—Bandwidth reserved for class 3

You can configure setup and holding priorities for an LSP, but the following restrictions apply:

- The combination of class and priority must be one of the configured traffic engineering classes. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit operation to fail.
- Automatic bandwidth allocation is not supported. If you configure automatic bandwidth allocation, the commit operation fails.
- LSPs configured with the **bandwidth** statement but without specifying a class type use the default class type **ct0**.
- For migration issues, see Internet draft draft-ietf-tewg-diff-te-proto-07.txt.

Configuring Policing for LSPs

Policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each LSP.

For information about how to configure a policer for an LSP, see [“Configuring Policers for LSPs” on page 121](#).

Configuring Fast Reroute for Traffic-Engineered LSPs

You can configure fast reroute for traffic engineered LSPs (LSPs carrying a single class of traffic). It is also possible to reserve bandwidth on the detour path for the class of traffic when fast reroute is enabled. The same class type number is used for both the traffic engineered LSP and its detour.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling DiffServ-aware traffic engineering and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

You can configure the amount of bandwidth to reserve for detours using either the **bandwidth** statement or the **bandwidth-percent** statement. You can only configure one these statements at a time. If you do not configure either the **bandwidth** statement or the **bandwidth-percent** statement, the default setting is to not reserve bandwidth for the detour path (the bandwidth guarantee will be lost if traffic is switched to the detour).

When you configure the **bandwidth** statement, you can specify the specific amount of bandwidth (in bits per second [bps]) you want to reserve for the detour path. For information, see [“Configuring Fast Reroute” on page 479](#).

The **bandwidth-percent** statement allows you to specify the bandwidth of the detour path as a percentage of the bandwidth configured for the protected path. For example, if you configure 100 millions bps of bandwidth for the protected path and configure 20 for the **bandwidth-percent** statement, the detour path will have 20 million bps of bandwidth reserved for its use.

To configure the percent of bandwidth used by the detour path based on the bandwidth of the protected path, include the **bandwidth-percent** statement:

```
bandwidth-percent percentage;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **fast-reroute**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **fast-reroute**]

Configuring Multiclass LSPs

IN THIS SECTION

- [Configuring Class of Service for the Interfaces | 1216](#)
- [Configuring the IGP | 1216](#)
- [Configuring Class-Type Bandwidth Constraints for Multiclass LSPs | 1216](#)
- [Configuring Policing for Multiclass LSPs | 1218](#)
- [Configuring Fast Reroute for Multiclass LSPs | 1218](#)

A multiclass LSP is an LSP configured to reserve bandwidth for multiple class types and also carries the traffic for these class types. The differentiated service behavior is determined by the EXP bits.

You must configure the Differentiated Services domain (see [“Configuring Routers for DiffServ-Aware Traffic Engineering” on page 1199](#)) before you can enable a multiclass LSP. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in a multiclass LSP configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the multiclass LSP to function properly.

NOTE: You must configure extended MAM as the bandwidth model when you configure multiclass LSPs. See [“Configuring the Bandwidth Model” on page 1201](#).

All the routers participating in a multiclass LSP must be Juniper Networks routers running Junos OS Release 6.2 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the Junos OS. However, the multiclass LSP cannot traverse these routers.

To enable multiclass LSPs, you need to configure the following:

Configuring Class of Service for the Interfaces

The existing class-of-service infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to consistently mark traffic are configured with the existing Junos OS CoS features.

NOTE: The Junos OS does not support ATM interfaces.

For information about how to configure CoS, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

Configuring the IGP

You can configure either IS-IS or OSPF. The IS-IS and OSPF configurations for routers supporting multiclass LSPs are standard. For information about how to configure these protocols, see the *Junos OS Routing Protocols Library*.

Configuring Class-Type Bandwidth Constraints for Multiclass LSPs

You configure a multiclass LSP by using the standard LSP configuration statements and procedures. To configure an LSP as a multiclass LSP, specify the class type bandwidth constraints by including the **bandwidth** statement:

```
bandwidth {  
  ct0 bps;  
  ct1 bps;  
  ct2 bps;  
  ct3 bps;  
}
```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for these statements.

The class type statements specify bandwidth (in bits per second) for the following classes:

- **ct0**—Bandwidth reserved for class 0
- **ct1**—Bandwidth reserved for class 1

- **ct2**—Bandwidth reserved for class 2
- **ct3**—Bandwidth reserved for class 3

For example, to configure 50 megabytes of bandwidth for class type 1 and 30 megabytes of bandwidth for class type 2, include the **bandwidth** statement as follows:

```
[edit protocols mpls]
label-switched-path traffic-class {
  bandwidth {
    ct1 50M;
    ct2 30M;
  }
}
```

You cannot configure a bandwidth for a class type and also configure a bandwidth at the **[edit protocols mpls label-switched-path *lsp-name* bandwidth]** hierarchy level. For example, the following configuration cannot be committed:

```
[edit protocols mpls]
label-switched-path traffic-class {
  bandwidth {
    20M;
    ct1 10M;
  }
}
```

You can configure setup and holding priorities for a multiclass LSP, but the following restrictions apply:

- The setup and holding priorities apply to all classes for which bandwidth is requested.
- The combination of class and priority must be one of the configured traffic engineering classes. The default traffic engineering class configuration results in multiclass LSPs that cannot preempt and cannot be preempted. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit operation to fail.
- Automatic bandwidth allocation is not supported for multiclass LSPs. If you configure automatic bandwidth allocation, the commit operation fails.
- LSPs configured with the **bandwidth** statement but without specifying a class type use the default class type **ct0**.

Configuring Policing for Multiclass LSPs

Policing allows you to control the amount of traffic forwarded through a particular multiclass LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each multiclass LSP. You can also enable automatic policing for multiclass LSPs.

For information about how to configure a policer for a multiclass LSP, see [“Configuring Policers for LSPs” on page 121](#) and [“Configuring Automatic Policers” on page 124](#).

Configuring Fast Reroute for Multiclass LSPs

You can enable fast reroute for multiclass LSPs. The bandwidth guarantees for the class types can be carried over to the detour path in case the primary path of the multiclass LSP fails. The same traffic class types configured for the primary multiclass LSP are also signaled for the detour LSP.

The bandwidth guarantee for the detour path is a percentage of the bandwidth configured for the class types of the primary path. For example, you configure a value of 50 percent for the detour path and the protected LSP carries traffic for class types CT0 through CT3. The detour path is signaled with the same class types (CT0 through CT3) but with 50 percent of the bandwidth configured for the protected LSP.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling DiffServ-aware traffic engineering, that all of the traffic class types needed are available, and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

The bandwidth percentage for fast reroute is signaled from the ingress router to the egress router. All of the intermediate devices must complete their own CSPF computations and signaling.

When you configure the **bandwidth-percent** statement, the detour path bandwidth is computed by multiplying by the bandwidth configured for the primary multiclass LSP. For information about how to configure the bandwidth for the multiclass LSP, see [“Configuring Traffic-Engineered LSPs” on page 1213](#).

To configure the percentage of bandwidth used by the detour path based on the bandwidth of the protected path, include the **bandwidth-percent** statement:

```
bandwidth-percent percentage;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **fast-reroute**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **fast-reroute**]



MPLS Transport Profile

Configuring Operation, Administration, and Maintenance (OAM) for MPLS | **1221**

Configuring MPLS Pseudowires | **1241**

Configuring Class-of-Service (CoS) for MPLS | **1311**

Configuring Generalized MPLS (GMPLS) | **1353**

Configuring Operation, Administration, and Maintenance (OAM) for MPLS

IN THIS CHAPTER

- [Configuring the MPLS Transport Profile for OAM | 1221](#)
- [Configuring OAM Ingress Policies for LDP | 1238](#)
- [Tracing MPLS and LSP Packets and Operations | 1239](#)

Configuring the MPLS Transport Profile for OAM

IN THIS SECTION

- [MPLS Transport Profile Overview | 1221](#)
- [Example: Configuring the MPLS Transport Profile for OAM | 1222](#)

MPLS Transport Profile Overview

RFC 5654, *Requirements of an MPLS Transport Profile*, describes the requirements for the MPLS Transport Profile (MPLS-TP) that extends capabilities for Operation, Administration, and Maintenance (OAM) when MPLS is used for transport services and transport network operations. These capabilities help in troubleshooting and maintenance of a pseudowire or label-switched path (LSP).

MPLS-TP mechanisms for OAM contain two main components:

- Generic Associated Channel Label (GAL)—A special label that enables an exception mechanism that informs the egress label-switching router (LSR) that a packet it receives on an LSP belongs to an associated control channel or the control plane.

- Generic Associated Channel Header (G-Ach)—A special header field that identifies the type of payload contained in the MPLS label-switched paths (LSPs). G-Ach has the same format as a pseudowire associated control channel header.

For more information about MPLS-TP, see RFC 5654, *Requirements of an MPLS Transport Profile*. For specific information about GAL and G-Ach, see RFC 5586, *MPLS Generic Associated Channel*.

The following capabilities are supported in the Junos OS implementation of MPLS-TP:

- MPLS-TP OAM can send and receive packets with GAL and G-Ach, without IP encapsulation.
- Two unidirectional RSVP LSPs between a pair of routers can be associated with each other to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages. A single Bidirectional Forwarding Detection (BFD) session is established for the associated bidirectional LSP.

Example: Configuring the MPLS Transport Profile for OAM

IN THIS SECTION

- [Requirements | 1222](#)
- [Overview | 1222](#)
- [Configuration | 1225](#)
- [Verification | 1236](#)

This example shows how to configure the MPLS Transport Profile (MPLS-TP) for sending and receiving of OAM GAL and G-Ach messages across a label-switched path (LSP).

Requirements

This example uses the following hardware and software components:

- Six devices that can be a combination of M Series, MX Series, and T Series routers
- Junos OS Release 12.1 or later running on the devices

Overview

Junos OS Release 12.1 and later support MPLS Transport Profile (MPLS-TP) Operation, Administration, and Maintenance (OAM) capabilities. MPLS-TP introduces new capabilities for OAM when MPLS is used for transport services and transport network operations. This includes configuring Generic Associated Channel Label (GAL) and Generic Associated Channel Header (G-Ach) for OAM messages.

This example shows how to configure MPLS-TP OAM capability to send and receive GAL and G-Ach OAM messages without IP encapsulation. In addition, it also shows how to associate two unidirectional RSVP

label-switched paths (LSPs) between a pair of routers to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages.

Junos OS Release 12.1 and later support the following MPLS-TP capabilities:

- MPLS-TP OAM capability and the infrastructure required for MPLS applications to send and receive packets with GAL and G-Ach, without IP encapsulation.
- LSP-ping and Bidirectional Forwarding Detection (BFD) applications to send and receive packets using GAL and G-Ach, without IP encapsulation on transport LSPs.
- The association of two unidirectional RSVP LSPs, between a pair of routers, with each other to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages. The associated bidirectional LSP model is supported only for associating the primary paths. A single BFD session is established for the associated bidirectional LSP.

Junos OS Release 12.1 and later does not support the following MPLS-TP capabilities:

- Point-to-multipoint RSVP LSPs and BGP LSPs
- Loss Measurement and Delay Measurement

You can enable GAL and G-Ach OAM operation using the following configuration statements:

- **mpls-tp-mode**—Include this statement at the **[edit protocols mpls oam]** hierarchy level to enable GAL and G-Ach OAM operation, without IP encapsulation, on all LSPs in the MPLS network.

```
[edit protocols mpls oam]
mpls-tp-mode;
```

Include this statement at the **[edit protocols mpls label-switched-path *lsp-name* oam]** hierarchy level to enable GAL and G-Ach OAM operation without IP encapsulation on a specific LSP in the network.

```
[edit protocols mpls label-switched-path lsp-name oam]
mpls-tp-mode;
```


NOTE: Starting with Junos OS Release 16.1, MPLS-TP supports two additional channel types for the default LSPING (0x0008) channel type under the `mpls-tp-mode` statement. These additional channel types provide on-demand connectivity verification (CV) with and without IP/UDP encapsulation.

- On-demand CV (0x0025)—This channel type is a new pseudowire channel type and is used for on-demand CV without IP/UDP encapsulation, where IP addressing is not available or non-IP encapsulation is preferred.
- IPv4 (0x0021)—This channel type uses the IP/UDP encapsulation and provides interoperability support with other vendor devices using IP addressing.

The GACH-TLV is used along with the default LSPING channel type. As per RFC 7026, GACH-TLV is deprecated for 0x0021 and 0x0025 channel types.

To configure a channel type for MPLS-TP, include the `lsping-channel-type channel-type` statement at the `[edit protocols mpls label-switched-path lsp-name oam mpls-tp-mode]` and `[edit protocols mpls oam mpls-tp-mode]` hierarchy levels.

- **associate-lsp *lsp-name* from *from-ip-address***—Include this statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level to configure associated bidirectional LSPs on the two ends of the LSP.

```
[edit protocols mpls label-switched-path lsp-name ]
associate-lsp lsp-name {
    from from-ip-address;
}
```

The **from *from-ip-address*** configuration for the LSP is optional. If omitted, it is derived from the **to** address of the ingress LSP configuration.

- **transit-lsp-association**—Include this statement at the `[edit protocols mpls]` hierarchy level to associate two LSPs at a transit router.


```
[edit protocols mpls]
transit-lsp-association transit-association-lsp-group-name {
    lsp-name-1 name-of-associated-lsp-1;
    from-1 address-of-associated-lsp-1;
    lsp-name-2 name-of-associated-lsp-2;
    from-2 address-of-associated-lsp-2;
}
```


The association of the LSPs in the transit nodes is useful for the return LSP path for TTL-expired LSP ping packets or traceroute.

In this example, R0 is the ingress router and R4 is the egress router. R1, R2, R3, and R5 are transit routers. The associated bidirectional LSP is established between the transit routers for sending and receiving the GAL and G-Ach OAM messages.

Figure 99 on page 1225 shows the topology used in this example.

Figure 99: MPLS-TP OAM Associated Bidirectional LSPs



Configuration

CLI Quick Configuration

NOTE: This example shows the configuration on all devices and shows step-by-step procedures for configuring the ingress router, R0, and transit router R1. Repeat the step-by-step procedure described for the ingress router, R0, on the egress router, R4. Repeat the step-by-step procedure for the transit router, R1, on the other transit routers, R2, R3, and R5. Be sure to modify the appropriate interface names, addresses, and other parameters appropriately.

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router R0

```
set interfaces ge-4/1/1 unit 0 family inet address 10.10.11.1/30
set interfaces ge-4/1/1 unit 0 family iso
set interfaces ge-4/1/1 unit 0 family inet6
set interfaces ge-4/1/1 unit 0 family mpls
set interfaces ge-5/0/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-5/0/0 unit 0 family iso
set interfaces ge-5/0/0 unit 0 family inet6
set interfaces ge-5/0/0 unit 0 family mpls
set protocols rsvp interface ge-5/0/0.0
set protocols rsvp interface ge-4/1/1.0
set protocols mpls label-switched-path r0-to-r4 to 10.255.8.86
set protocols mpls label-switched-path r0-to-r4 oam mpls-tp-mode
set protocols mpls label-switched-path r0-to-r4 associate-lsp r4-to-r0 from 10.255.8.86
set protocols mpls interface ge-5/0/0.0
set protocols mpls interface ge-4/1/1.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-5/0/0.0
set protocols ospf area 0.0.0.0 interface ge-4/1/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Router R1

```
set interfaces ge-0/0/5 unit 0 family inet address 10.10.10.2/30
set interfaces ge-0/0/5 unit 0 family iso
set interfaces ge-0/0/5 unit 0 family inet6
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/2/2 unit 0 family inet address 10.10.12.2/30
set interfaces ge-0/2/2 unit 0 family iso
set interfaces ge-0/2/2 unit 0 family inet6
set interfaces ge-0/2/2 unit 0 family mpls
set interfaces ge-1/0/2 unit 0 family inet address 10.10.13.2/30
set interfaces ge-1/0/2 unit 0 family iso
set interfaces ge-1/0/2 unit 0 family inet6
set interfaces ge-1/0/2 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 family inet address 10.10.11.2/30
set interfaces ge-2/0/2 unit 0 family iso
set interfaces ge-2/0/2 unit 0 family inet6
set interfaces ge-2/0/2 unit 0 family mpls
```



```

set protocols rsvp interface ge-0/2/2.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-1/0/2.0
set protocols rsvp interface ge-2/0/2.0
set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86
set protocols mpls interface ge-0/0/5.0
set protocols mpls interface ge-2/0/2.0
set protocols mpls interface ge-1/0/2.0
set protocols mpls interface ge-0/2/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/2/2.0 metric 100
set protocols ospf area 0.0.0.0 interface ge-1/0/2.0
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Router R2

```

set interfaces ge-0/2/3 unit 0 family inet address 10.10.13.1/30
set interfaces ge-0/2/3 unit 0 family iso
set interfaces ge-0/2/3 unit 0 family inet6
set interfaces ge-0/2/3 unit 0 family mpls
set interfaces ge-1/3/2 unit 0 family inet address 10.10.14.1/30
set interfaces ge-1/3/2 unit 0 family iso
set interfaces ge-1/3/2 unit 0 family inet6
set interfaces ge-1/3/2 unit 0 family mpls
set interfaces ge-1/3/4 unit 0 family inet address 10.10.15.1/30
set interfaces ge-1/3/4 unit 0 family iso
set interfaces ge-1/3/4 unit 0 family inet6
set interfaces ge-1/3/4 unit 0 family mpls
set protocols rsvp interface ge-0/2/3.0
set protocols rsvp interface ge-1/3/2.0
set protocols rsvp interface ge-1/3/4.0
set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86

```



```

set protocols mpls interface ge-0/2/3.0
set protocols mpls interface ge-1/3/2.0
set protocols mpls interface ge-1/3/4.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/2/3.0
set protocols ospf area 0.0.0.0 interface ge-1/3/2.0
set protocols ospf area 0.0.0.0 interface ge-1/3/4.0 metric 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Router R3

```

set interfaces ge-1/2/1 unit 0 family inet address 10.10.16.2/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family inet6
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-2/0/7 unit 0 family inet address 10.10.17.2/30
set interfaces ge-2/0/7 unit 0 family iso
set interfaces ge-2/0/7 unit 0 family inet6
set interfaces ge-2/0/7 unit 0 family mpls
set interfaces ge-2/2/0 unit 0 family inet address 10.10.14.2/30
set interfaces ge-2/2/0 unit 0 family iso
set interfaces ge-2/2/0 unit 0 family inet6
set interfaces ge-2/2/0 unit 0 family mpls
set protocols rsvp interface ge-2/2/0.0
set protocols rsvp interface ge-1/2/1.0
set protocols rsvp interface ge-2/0/7.0
set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86
set protocols mpls interface ge-2/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface ge-2/0/7.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/2/0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ospf area 0.0.0.0 interface ge-2/0/7.0 metric 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Router R4


```

set interfaces ge-0/0/3 unit 0 family inet address 10.10.16.1/30
set interfaces ge-0/0/3 unit 0 family iso
set interfaces ge-0/0/3 unit 0 family inet6
set interfaces ge-0/0/3 unit 0 family mpls
set protocols rsvp interface ge-0/0/3.0
set protocols mpls label-switched-path r4-to-r0 to 10.255.8.207
set protocols mpls label-switched-path r4-to-r0 oam mpls-tp-mode
set protocols mpls label-switched-path r4-to-r0 associate-lsp r0-to-r4 from 10.255.8.207
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Router R5

```

set interfaces ge-1/2/0 unit 0 family inet address 10.10.15.2/30
set interfaces ge-1/2/0 unit 0 family iso
set interfaces ge-1/2/0 unit 0 family inet6
set interfaces ge-1/2/0 unit 0 family mpls
set interfaces ge-2/0/0 unit 0 family inet address 10.10.12.1/30
set interfaces ge-2/0/0 unit 0 family iso
set interfaces ge-2/0/0 unit 0 family inet6
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-4/0/7 unit 0 family inet address 10.10.17.1/30
set interfaces ge-4/0/7 unit 0 family iso
set interfaces ge-4/0/7 unit 0 family inet6
set interfaces ge-4/0/7 unit 0 family mpls
set protocols rsvp interface ge-2/0/0.0
set protocols rsvp interface ge-1/2/0.0
set protocols rsvp interface ge-4/0/7.0
set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86
set protocols mpls interface ge-2/0/0.0
set protocols mpls interface ge-1/2/0.0
set protocols mpls interface ge-4/0/7.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 100
set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 metric 100

```



```
set protocols ospf area 0.0.0.0 interface ge-4/0/7.0 metric 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

Configuring Device R0

Step-by-Step Procedure

To configure the ingress router, R0:

1. Configure the interfaces.

```
[edit interfaces]
user@R0# set ge-4/1/1 unit 0 family inet address 10.10.11.1/30
user@R0# set ge-4/1/1 unit 0 family iso
user@R0# set ge-4/1/1 unit 0 family inet6
user@R0# set ge-4/1/1 unit 0 family mpls
user@R0# set ge-5/0/0 unit 0 family inet address 10.10.10.1/30
user@R0# set ge-5/0/0 unit 0 family iso
user@R0# set ge-5/0/0 unit 0 family inet6
user@R0# set ge-5/0/0 unit 0 family mpls
```

2. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@R0# set interface ge-5/0/0.0
user@R0# set interface ge-4/1/1.0
```

3. Configure an interior gateway protocol, such as OSPF.

```
[edit protocols ospf]
user@R0# set traffic-engineering
user@R0# set area 0.0.0.0 interface ge-5/0/0.0
user@R0# set area 0.0.0.0 interface ge-4/1/1.0
user@R0# set area 0.0.0.0 interface lo0.0 passive
```

4. Configure a signaling protocol, such as RSVP.

```
[edit protocols rsvp]
user@R0# set interface ge-5/0/0.0
user@R0# set interface ge-4/1/1.0
```


5. Configure the LSP.

```
[edit protocols mpls]
user@R0# set label-switched-path r0-to-r4 to 10.255.8.86
```

6. Enable GAL and G-Ach OAM operation without IP encapsulation on the LSPs.

```
[edit protocols mpls]
user@R0# set label-switched-path r0-to-r4 oam mpls-tp-mode
```

7. Configure associated bidirectional LSPs on the two ends of the LSP.

```
[edit protocols mpls]
user@R0# set label-switched-path r0-to-r4 associate-lsp to-r0 from 10.255.8.86
```

8. After you are done configuring the device, commit the configuration.

```
[edit]
user@R0# commit
```

Results

Confirm your configuration by issuing the **show interfaces** and **show protocols** commands.

```
user@R0# show interfaces
ge-4/1/1 {
  unit 0 {
    family inet {
      address 10.10.11.1/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-5/0/0 {
  unit 0 {
    family inet {
      address 10.10.10.1/30;
    }
    family iso;
```



```

    family inet6;
    family mpls;
  }
}

```

```

user@R0# show protocols
rsvp {
  interface ge-5/0/0.0;
  interface ge-4/1/1.0;
}
mpls {
  label-switched-path r0-to-r4 {
    to 10.255.8.86;
    oam mpls-tp-mode;
    associate-lsp r4-to-r0 {
      from 10.255.8.86;
    }
  }
  interface ge-4/1/1.0;
  interface ge-5/0/0.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-5/0/0.0;
    interface ge-4/1/1.0;
    interface lo0.0 {
      passive;
    }
  }
}
}

```

Configuring Device R1

Step-by-Step Procedure

To configure the transit router, R1:

1. Configure the interfaces.

```

[edit interfaces]
user@R1# set ge-0/0/5 unit 0 family inet address 10.10.10.2/30
user@R1# set ge-0/0/5 unit 0 family iso
user@R1# set ge-0/0/5 unit 0 family inet6
user@R1# set ge-0/0/5 unit 0 family mpls

```



```

user@R1# set ge-0/2/2 unit 0 family inet address 10.10.12.2/30
user@R1# set ge-0/2/2 unit 0 family iso
user@R1# set ge-0/2/2 unit 0 family inet6
user@R1# set ge-0/2/2 unit 0 family mpls
user@R1# set ge-2/0/2 unit 0 family inet address 10.10.11.2/30
user@R1# set ge-2/0/2 unit 0 family iso
user@R1# set ge-2/0/2 unit 0 family inet6
user@R1# set ge-2/0/2 unit 0 family mpls
user@R1# set ge-1/0/2 unit 0 family inet address 10.10.13.2/30
user@R1# set ge-1/0/2 unit 0 family iso
user@R1# set ge-1/0/2 unit 0 family inet6
user@R1# set ge-1/0/2 unit 0 family mpls

```

2. Configure MPLS on the interfaces.

```

[edit protocols mpls]
user@R1# set interface ge-0/0/5.0
user@R1# set interface ge-2/0/2.0
user@R1# set interface ge-1/0/2.0
user@R1# set interface ge-0/2/2.0

```

3. Configure an interior gateway protocol, such as OSPF.

```

[edit protocols ospf]
user@R1# set traffic-engineering
user@R1# set area 0.0.0.0 interface ge-0/0/5.0
user@R1# set area 0.0.0.0 interface ge-2/0/2.0
user@R1# set area 0.0.0.0 interface ge-1/0/2.0
user@R1# set area 0.0.0.0 interface ge-0/2/2.0 metric 100
user@R1# set area 0.0.0.0 interface lo0.0 passive

```

4. Configure a signaling protocol, such as RSVP.

```

[edit protocols rsvp]
user@R1# set interface ge-0/0/5.0
user@R1# set interface ge-2/0/2.0
user@R1# set interface ge-1/0/2.0
user@R1# set interface ge-0/2/2.0

```

5. Configure the association of the two LSPs on the transit router.


```
[edit protocols mpls]
```

```
user@R1# set transit-lsp-association trace1 lsp-name-1 r0-to-r4
```

```
user@R1# set transit-lsp-association trace1 from-1 10.255.8.207
```

```
user@R1# set transit-lsp-association trace1 lsp-name-2 r4-to-r0
```

```
user@R1# set transit-lsp-association trace1 from-2 10.255.8.86
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
```

```
user@R1# commit
```

Results

Confirm your configuration by issuing the **show interfaces** and **show protocols** commands.

```
user@R1# show interfaces
ge-0/0/5 {
  unit 0 {
    family inet {
      address 10.10.10.2/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-0/2/2 {
  unit 0 {
    family inet {
      address 10.10.12.2/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-2/0/2 {
  unit 0 {
    family inet {
      address 10.10.11.2/30;
    }
    family iso;
    family inet6;
```



```

        family mpls;
    }
}
ge-1/0/2 {
    unit 0 {
        family inet {
            address 10.10.13.2/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}

```

```

user@R1# show protocols
rsvp {
    interface ge-0/0/5.0;
    interface ge-2/0/2.0;
    interface ge-1/0/2.0;
    interface ge-0/2/2.0;
}
mpls {
    transit-lsp-association trace1 {
        lsp-name-1 r0-to-r4;
        from-1 10.255.8.207;
        lsp-name-2 r4-to-r0;
        from-2 10.255.8.86;
    }
    interface ge-0/0/5.0;
    interface ge-2/0/2.0;
    interface ge-1/0/2.0;
    interface ge-0/2/2.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/0/5.0;
        interface ge-1/0/2.0;
        interface ge-2/0/2.0;
        interface ge-0/2/2.0 {
            metric 100;
        }
        interface lo0.0 {
            passive;

```



```

    }
  }
}

```

Verification

Confirm that the configuration is working properly.

Verifying Associated Bidirectional LSPs

Purpose

Verify that the associated bidirectional LSP configuration is working properly.

Action

```
user@host> show mpls lsp
```

```

Ingress LSP: 1 sessions
To          From          State Rt P    ActivePath  LSPname
10.10.11.1  10.255.8.86      Up    0 *           r0-to-r4 Assoc-Bidir
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.10.16.1  10.255.8.207  Up    0  1 FF         3          r4-to-r0 Assoc-Bidir
Total 2 displayed, Up 2, Down 0

Transit LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.10.10.2  10.255.8.168  Up    1  1 FF      301264      3 r0-to-r4 Assoc-Bidir
Total 3 displayed, Up 3, Down 0

```

```
user@host> show mpls lsp detail
```

```

Ingress LSP: 1 sessions

10.10.11.1
  From: 10.255.8.86, State: Up, ActiveRoute: 0, LSPname: r0-to-r4
  Associated Bidirectional
  Associated LSP: r0-to-r4, 10.255.8.86
  ActivePath: (primary)
  LSPTYPE: Static Configured
  LoadBalance: Random

```



```

    Encoding type: Packet, Switching type: PSC-1, GPID: Unknown
    *Primary                               State: Up

Egress LSP: 1 sessions

10.255.102.29
  From: 10.255.102.172, LSPstate: Up, ActiveRoute: 0
    LSPname: r4-to-r0, LSPpath: Primary
    Associated Bidirectional
    Associated LSP: 10.10.16.1, to-r0>
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: -
    Resv style: 1 FF, Label in: 3, Label out: -
    Time left: 144, Since: Fri Jun 17 21:41:05 2011
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 6 receiver 14468 protocol 0
    PATH rcvfrom: 10.10.13.1 (ge-2/0/0.0) 84 pkts
    Adspec: received MTU 1500
    PATH sentto: localclient
    RESV rcvfrom: localclient
    Record route: 10.10.14.2 10.10.13.1 <self>

Transit LSP: 1 sessions

10.255.102.30
  From: 10.255.102.172, LSPstate: Up, ActiveRoute: 1
    LSPname: to_airstream, LSPpath: Primary
    Associated Bidirectional
    Associated LSP: r0-to-r4, 10.255.8.168
    Suggested label received: -, Suggested label
    Recovery label received: -, Recovery label sent: 3
    Resv style: 1 FF, Label in: 301264, Label out: 3
    Time left: 132, Since: Fri Jun 17 21:40:56 2011
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 28 receiver 14465 protocol 0
    PATH rcvfrom: 10.10.13.1 (ge-2/0/0.0) 84 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.10.10.1 (ge-3/0/0.0) 84 pkts
    RESV rcvfrom: 10.10.10.1 (ge-3/0/0.0) 84 pkts
    Explct route: 10.10.10.1
    Record route: 10.10.16.1 10.10.15.2 10.10.13.1 <self> 10.10.10.1

```

user@host> **show mpls lsp bidirectional**


```

Ingress LSP: 1 session
To          From          State Rt P      ActivePath      LSPname
10.255.8.86  10.255.8.207  Up    0 *
  Assoc-Bidir
Total 1 displayed, Up 1, Down 0
Aug 28 06:56:26 [TRACE] [R0 coleman re0]
Egress LSP: 1 session
To          From          State  Rt Style Labelin Labelout LSPname
10.255.8.207  10.255.8.86  Up      0  1 FF      3      - to-r0
  Assoc-Bidir
Total 1 displayed, Up 1, Down 0
Aug 28 06:56:26 [TRACE] [R0 coleman re0]
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The output of the **show mpls lsp**, **show mpls detail**, and **show mpls bidirectional** commands displays the details of the associated bidirectional LSPs and the LSP association information.

SEE ALSO

[MPLS Transport Profile Overview | 1221](#)

[associate-lsp | 2558](#)

[mpls-tp-mode | 2690](#)

[transit-lsp-association | 2823](#)

Configuring OAM Ingress Policies for LDP

Using the **ingress-policy** statement, you can configure an Operation, Administration, and Management (OAM) policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under **[edit protocols ldp oam bfd-liveness-detection]** are applied.

You configure the OAM ingress policy at the **[edit policy-options]** hierarchy level. To configure an OAM ingress policy, include the **ingress-policy** statement:

```
ingress-policy ingress-policy-name;
```


You can configure this statement at the following hierarchy levels:

- [edit protocols ldp oam]
- [edit logical-systems *logical-system-name* protocols ldp oam]

NOTE: ACX Series routers do not support [edit logical-systems] hierarchy level.

Tracing MPLS and LSP Packets and Operations

To trace MPLS and LSP packets and operations, include the **traceoptions** statement:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can specify the following MPLS-specific flags in the MPLS **traceoptions** statement:

- **all**—Trace all operations.
- **connection**—Trace all circuit cross-connect (CCC) activity.
- **connection-detail**—Trace detailed CCC activity.
- **cspf**—Trace CSPF computations.
- **cspf-link**—Trace links visited during CSPF computations.
- **cspf-node**—Trace nodes visited during CSPF computations.
- **error**—Trace MPLS error conditions.
- **graceful-restart**—Trace MPLS graceful restart events.
- **lsping**—Trace LSP ping packets and return codes.
- **nsr-synchronization**—Trace nonstop routing (NSR) synchronization events.
- **nsr-synchronization-detail**—Trace NSR synchronization events in detail.
- **state**—Trace all LSP state transitions.
- **static**—Trace static label-switched path.

When you configure trace options to track an MPLS LSP using the **cspf** option, the CSPF log displays information about the MPLS LSP using the term “generalized MPLS” (GMPLS). For example, a message in the CSPF log might state that the “link passes GMPLS constraints”. Generalized MPLS (GMPLS) is a superset of MPLS, so this message is normal and does not affect proper MPLS LSP operation.

For general information about tracing and global tracing options, see the *Junos OS Routing Protocols Library*.

Configuring MPLS Pseudowires

IN THIS CHAPTER

- [Ethernet Pseudowire Overview | 1241](#)
- [Example: Ethernet Pseudowire Base Configuration | 1242](#)
- [Pseudowire Overview for ACX Series Universal Metro Routers | 1246](#)
- [Understanding Multisegment Pseudowire for FEC 129 | 1248](#)
- [Example: Configuring a Multisegment Pseudowire | 1253](#)
- [MPLS Stitching For Virtual Machine Connection | 1299](#)
- [TDM Pseudowires Overview | 1302](#)
- [Example: TDM Pseudowire Base Configuration | 1303](#)
- [Configuring Load Balancing for Ethernet Pseudowires | 1308](#)
- [Configuring Load Balancing Based on MAC Addresses | 1309](#)

Ethernet Pseudowire Overview

Starting in Junos OS Release 14.1X53 and Junos OS Release 16.1, 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.

Release History Table

Release	Description
14.1X53	Starting in Junos OS Release 14.1X53 and Junos OS Release 16.1, 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.

RELATED DOCUMENTATION

Configuring IEEE 802.3ah OAM Link-Fault Management

[Pseudowire Overview for ACX Series Universal Metro Routers](#) | [1246](#)

[TDM Pseudowires Overview](#) | [1302](#)

ATM Pseudowire Overview

Example: Ethernet Pseudowire Base Configuration

IN THIS SECTION

- [Requirements](#) | [1242](#)
- [Overview of an Ethernet Pseudowire Base Configuration](#) | [1243](#)
- [Configuring an Ethernet Pseudowire](#) | [1243](#)

Requirements

The following is a list of the hardware and software requirements for this configuration.

- One ACX Series router

- Junos OS Release 12.2 or later

Overview of an Ethernet Pseudowire Base Configuration

The configuration shown here is the base configuration of an Ethernet pseudowire with Ethernet cross-connect for physical interface encapsulation on an ACX Series router. This configuration is for one provider edge router. To complete the configuration of an Ethernet pseudowire, you need to repeat this configuration on an other provider edge router in the Multiprotocol Label Switched (MPLS) network.

Configuring an Ethernet Pseudowire

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in 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:

```
set interfaces ge-0/1/1 encapsulation ethernet-ccc
set interfaces ge-0/1/1 unit 0
set interfaces ge-0/2/0 unit 0 family inet address 20.1.1.2/24
set interfaces ge-0/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 70.1.1.1/32
set protocols rsvp interface ge-0/2/0.0
set protocols mpls no-cspf
set protocols mpls label-switched-path PE1-to-PE2 to 40.1.1.1
set protocols mpls interface ge-0/2/0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-0/2/0.0
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 40.1.1.1 interface ge-0/1/1.0 virtual-circuit-id
1
```

NOTE: To configure an Ethernet pseudowire with 802.1Q tagging for cross-connect logical interface encapsulation, include the **vlan-ccc** statement at the **[edit interfaces ge-0/1/1 unit 0 encapsulation]** hierarchy level instead of the **ethernet-ccc** statement shown in this example.

Step-by-Step Procedure

1. Create two Gigabit Ethernet interfaces, set the encapsulation mode on one interface and MPLS on the other interface. Create the loopback (**lo0**) interface:

```
[edit]
user@host# edit interfaces
[edit interfaces]
user@host# set ge-0/1/1 encapsulation ethernet-ccc
user@host# set ge-0/1/1 unit 0
user@host# set ge-0/2/0 unit 0 family inet address 20.1.1.2/24
user@host# set ge-0/2/0 unit 0 family mpls
user@host# set lo0 unit 0 family inet address 70.1.1.1/32
```

2. Enable the MPLS and RSVP protocols on the interface configured with MPLS—**ge-0/2/0.0**:

```
[edit]
user@host# edit protocols
[edit protocols]
user@host# set rsvp interface ge-0/2/0.0
user@host# set mpls interface ge-0/2/0.0
```

3. Configure LDP. If you configure RSVP for a pseudowire, you must also configure LDP:

```
[edit protocols]
user@host# set protocols ldp interface ge-0/2/0.0
user@host# set protocols ldp interface lo0.0
```

4. Configure a point-to-point label-switched path (LSP) and disable constrained-path LSP computation:

```
[edit protocols]
user@host# set mpls label-switched-path PE1-to-PE2 to 40.1.1.1
user@host# set mpls no-cspf
```

5. Configure OSPF and enable traffic engineering on the MPLS interface—**ge-0/2/0.0**, and on the loopback (**lo0**) interface:

```
[edit protocols]
user@host# set ospf traffic-engineering
user@host# set ospf area 0.0.0.0 interface ge-0/2/0.0
user@host# set ospf area 0.0.0.0 interface lo0.0 passive
```


6. Uniquely identify a Layer 2 circuit for the Ethernet pseudowire:

```
[edit protocols]
user@host# set l2circuit neighbor 40.1.1.1 interface ge-0/1/1.0 virtual-circuit-id 1
```

Results

```
[edit]
user@host# show
interfaces {
  ge-0/1/1 {
    encapsulation ethernet-ccc;
    unit 0;
  }
  ge-0/2/0 {
    unit 0 {
      family inet {
        address 20.1.1.2/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 70.1.1.1/32;
      }
    }
  }
}
protocols {
  rsvp {
    interface ge-0/2/0.0;
  }
  mpls {
    no-cspf;
    label-switched-path PE1-to-PE2 {
      to 40.1.1.1;
    }
    interface ge-0/2/0.0;
  }
  ospf {
    traffic-engineering;
```



```

        area 0.0.0.0 {
            interface ge-0/2/0.0;
            interface lo0.0 {
                passive;
            }
        }
    }
    ldp {
        interface ge-0/2/0.0;
        interface lo0.0;
    }
    l2circuit {
        neighbor 40.1.1.1 {
            interface ge-0/1/1.0 {
                virtual-circuit-id 1;
            }
        }
    }
}

```

RELATED DOCUMENTATION

[Pseudowire Overview for ACX Series Universal Metro Routers](#) | 1246

[Ethernet Pseudowire Overview](#) | 1241

Pseudowire Overview for ACX Series Universal Metro Routers

A pseudowire is a Layer 2 circuit or service, which emulates the essential attributes of a telecommunications service—such as a T1 line, over an MPLS packet-switched network. The pseudowire is intended to provide only the minimum necessary functionality to emulate the wire with the required degree of faithfulness for the given service definition. On the ACX Series routers, Ethernet, Asynchronous Transfer Mode (ATM), and time-division multiplexing (TDM) pseudowires are supported. The following pseudowire features are supported:

- Pseudowire transport service carrying Layer 1 and Layer 2 information over an IP and MPLS network infrastructure. Only similar end points are supported on the ACX Series—for example, T1 to T1, ATM to ATM, and Ethernet to Ethernet.
- Redundant pseudowires backup connections between PE routers and CE devices, maintaining Layer 2 circuits and services after certain types of failures. Pseudowire redundancy improves the reliability of

certain types of networks (metro for example) where a single point of failure could interrupt service for multiple customers. The following pseudowire redundancy features are supported:

- Maintenance of Layer 2 circuit services after certain types of failures with a standby pseudowire, which backs up the connection between PE routers and CE devices.
- In case of failure, a protect interface, which backs up the primary interface. Network traffic uses the primary interface only so long as the primary interface functions. If the primary interface fails, traffic is switched to the protect interface.
- Hot and cold standby enabling swift cut over to the backup or standby pseudowire.
- Ethernet connectivity fault management (CFM), which can be used to monitor the physical link between two routers. The following major features of CFM for Ethernet pseudowires only are supported:
 - 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

Layer 2 Circuits Overview

Redundant Pseudowires for Layer 2 Circuits and VPLS

Configuring a MEP to Generate and Respond to CFM Protocol Messages

IEEE 802.1ag OAM Connectivity Fault Management Overview

Configuring a CFM Action Profile to Specify CFM Actions for CFM Events

Configuring Interfaces for Layer 2 Circuits

[TDM Pseudowires Overview | 1302](#)

ATM Pseudowire Overview

[Ethernet Pseudowire Overview | 1241](#)

Understanding Multisegment Pseudowire for FEC 129

IN THIS SECTION

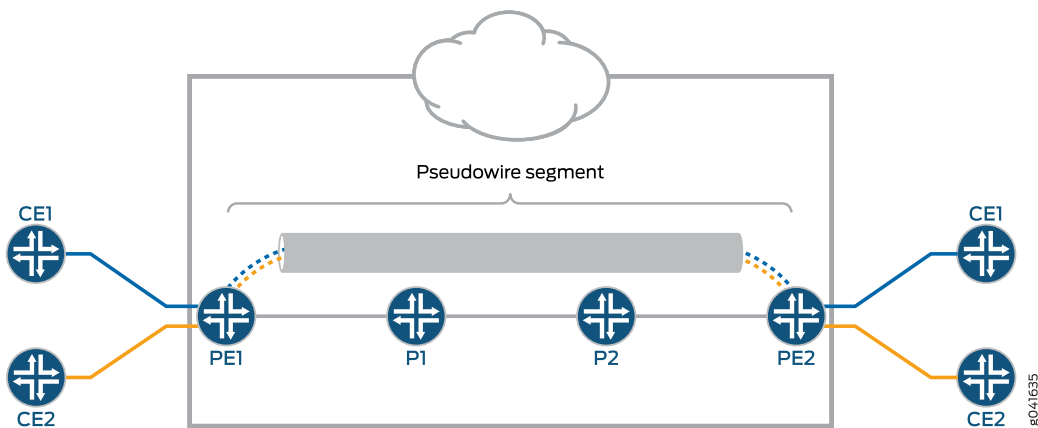
- Understanding Multisegment Pseudowire | 1248
- Using FEC 129 for Multisegment Pseudowire | 1249
- Establishing a Multisegment Pseudowire Overview | 1250
- Pseudowire Status Support for Multisegment Pseudowire | 1250
- Pseudowire TLV Support for MS-PW | 1251
- Supported and Unsupported Features | 1252

Understanding Multisegment Pseudowire

A pseudowire is a Layer 2 circuit or service that emulates the essential attributes of a telecommunications service, such as a T1 line, over an MPLS packet-switched network (PSN). The pseudowire is intended to provide only the minimum necessary functionality to emulate the wire with the required resiliency requirements for the given service definition.

When a pseudowire originates and terminates on the edge of the same PSN, the pseudowire label is unchanged between the originating and terminating provider edge (T-PE) devices. This is called a single-segment pseudowire (SS-PW). [Figure 100 on page 1248](#) illustrates an SS-PW established between two PE routers. The pseudowires between the PE1 and PE2 routers are located within the same autonomous system (AS).

Figure 100: L2VPN Pseudowire

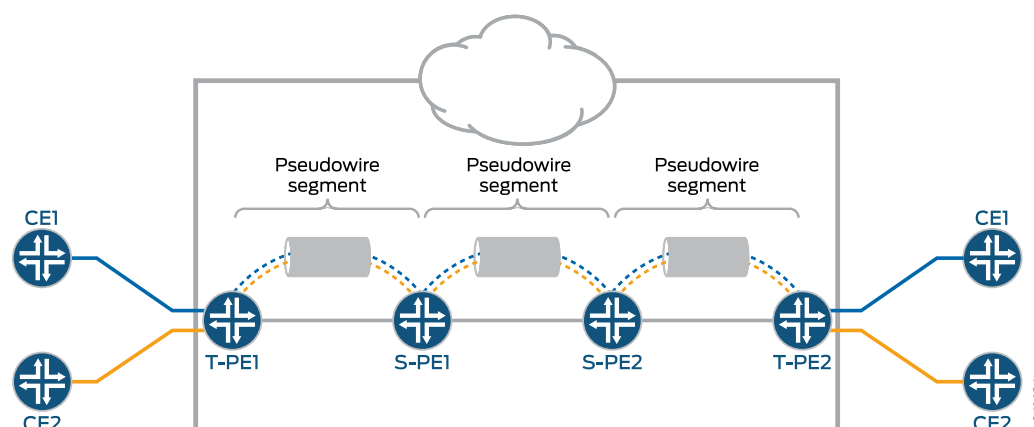


In cases where it is impossible to establish a single pseudowire from a local to a remote PE, either because it is unfeasible or undesirable to establish a single control plane between the two PEs, a multisegment pseudowire (MS-PW) is used.

An MS-PW is a set of two or more contiguous SS-PWs that are made to function as a single point-to-point pseudowire. It is also known as switched pseudowire. MS-PWs can go across different regions or network domains. A region can be considered as an interior gateway protocol (IGP) area or a BGP autonomous system that belongs to the same or different administrative domain. An MS-PW spans multiple cores or ASs of the same or different carrier networks. A Layer 2 VPN MS-PW can include up to 254 pseudowire segments.

Figure 101 on page 1249 illustrates a set of two or more pseudowire segments that function as a single pseudowire. The end routers are called terminating PE (T-PE) routers, and the switching routers are called switching PE (S-PE) routers. The S-PE router terminates the tunnels of the preceding and succeeding pseudowire segments in an MS-PW. The S-PE router can switch the control and data planes of the preceding and succeeding pseudowire segments of the MS-PW. An MS-PW is declared to be up when all the single-segment pseudowires are up.

Figure 101: Multisegment Pseudowire



Using FEC 129 for Multisegment Pseudowire

Currently, there are two types of attachment circuit identifiers (AIs) defined under FEC 129:

- Type 1 AI
- Type 2 AI

The support of an MS-PW for FEC 129 uses type 2 AI. A type 2 AI is globally unique by definition of RFC 5003.

Single-segment pseudowires (SS-PWs) using FEC 129 on an MPLS PSN can use both type 1 and type 2 All. For an MS-PW using FEC 129, a pseudowire itself is identified as a pair of endpoints. This requires that the pseudowire endpoints be uniquely identified.

In the case of a dynamically placed MS-PW, there is a requirement for the identifiers of attachment circuits to be globally unique, for the purposes of reachability and manageability of the pseudowire. Thus, individual globally unique addresses are allocated to all the attachment circuits and S-PEs that make up an MS-PW.

Type 2 All is composed of three fields:

- Global_ID—Global identification, which is usually the AS number.
- Prefix—IPv4 address, which is usually the router ID.
- AC_ID—Local attachment circuit, which is a user-configurable value.

Since type 2 All already contains the T-PE's IP address and it is globally unique, from the FEC 129 pseudowire signaling point of view, the combination (AGI, SAll, TAll) uniquely identifies an MS-PW across all interconnected pseudowire domains.

Establishing a Multisegment Pseudowire Overview

An MS-PW is established by dynamically and automatically selecting the predefined S-PEs and placing the MS-PW between two T-PE devices.

When S-PEs are dynamically selected, each S-PE is automatically discovered and selected using the BGP autodiscovery feature, without the requirement of provisioning the FEC 129 pseudowire-related information on all the S-PEs. BGP is used to propagate pseudowire address information throughout the PSN.

Since there is no manual provisioning of FEC 129 pseudowire information on the S-PEs, the Attachment Group Identifier (AGI) and Attachment Individual Identifier (AII) are reused automatically, and choosing the same set of S-PEs for the pseudowire in both the forwarding and reverse direction is achieved through the active and passive role of each T-PE device.

- Active—The T-PE initiates an LDP label mapping message.
- Passive—The T-PE does not initiate an LDP label mapping message until it receives a label mapping message initiated by the active T-PE. The passive T-PE sends its label mapping message to the same S-PE from where it received the label mapping message originated from its active T-PE. This ensures that the same set of S-PEs are used in the reverse direction.

Pseudowire Status Support for Multisegment Pseudowire

Pseudowire Status Behavior on T-PE

The following pseudowire status messages are relevant on the T-PE:

- 0x00000010—Local PSN-facing pseudowire (egress) transmit fault.

- 0x00000001—Generic nonforwarding fault code. This is set as the local fault code. The local fault code is set at the local T-PE, and LDP sends a pseudowire status TLV message with the same fault code to the remote T-PE.
- Fault codes are bit-wise OR'ed and stored as remote pseudowire status codes.

Pseudowire Status Behavior on S-PE

The S-PE initiates the pseudowire status messages that indicate the pseudowire faults. The SP-PE in the pseudowire notification message hints where the fault was originated.

- When a local fault is detected by the S-PE, a pseudowire status message is sent in both directions along the pseudowire. Since there are no attachment circuits on an S-PE, only the following status messages are relevant:
 - 0x00000008—Local PSN-facing pseudowire (ingress) receive fault.
 - 0x00000010—Local PSN-facing pseudowire (egress) transmit fault.
- To indicate which SS-PW is at fault, an LDP SP-PE TLV is attached with the pseudowire status code in the LDP notification message. The pseudowire status is passed along from one pseudowire to another unchanged by the control plane switching function.
- If an S-PE initiates a pseudowire status notification message with one particular pseudowire status bit, then for the pseudowire status code an S-PE receives, the same bit is processed locally and not forwarded until the S-PE's original status error is cleared.
- An S-PE keeps only two pseudowire status codes for each SS-PW it is involved in – local pseudowire status code and remote pseudowire status code. The value of the remote pseudowire status code is the result of logic or operation of the pseudowire status codes in the chain of SS-PWs preceding this segment. This status code is incrementally updated by each S-PE upon receipt and communicated to the next S-PE. The local pseudowire status is generated locally based on its local pseudowire status.
- Only transmit fault is detected at the SP-PE. When there is no MPLS LSP to reach the next segment, a local transmit fault is detected. The transmit fault is sent to the next downstream segment, and the receive fault is sent to the upstream segment.
- Remote failures received on an S-PE are just passed along the MS-PW unchanged. Local failures are sent to both segments of the pseudowire that the S-PE is involved in.

Pseudowire TLV Support for MS-PW

MS-PW provides the following support for the LDP SP-PE TLV [RFC 6073]:

- The LDP SP-PE TLVs for an MS-PW include:
 - Local IP address
 - Remote IP address

- An SP-PE adds the LDP SP-PE TLV to the label mapping message. Each SP-PE appends the local LDP SP-PE TLV to the SP-PE list it received from the other segment.
- The pseudowire status notification message includes the LDP SP-PE TLV when the notification is generated at the SP-PE.

Supported and Unsupported Features

Junos OS supports the following features with MS-PW:

- MPLS PSN for each SS-PW that builds up the MS-PW.
- The same pseudowire encapsulation for each SS-PW in an MS-PW – Ethernet or VLAN-CCC.
- The generalized PWid FEC with T-LDP as an end-to-end pseudowire signaling protocol to set up each SS-PW.
- MP-BGP to autodiscover the two endpoint PEs for each SS-PW associated with the MS-PW.
- Standard MPLS operation to stitch two side-by-side SS-PWs to form an MS-PW.
- Automatic discovery of S-PE so that the MS-PW can be dynamically placed.
- Minimum provisioning of S-PE.
- Operation, administration, and maintenance (OAM) mechanisms, including end-to-end MPLS ping or end-to-any-S-PE MPLS ping, MPLS path trace, end-to-end VCCV, and Bidirectional Forwarding Detection (BFD).
- Pseudowire swithing point (SP) PE TLV for the MS-PW.
- Composite next hop on MS-PW.
- Pseudowire status TLV for MS-PW.

Junos OS does not support the following MS-PW functionality:

- Mix of LDP FEC 128 and LDP FEC 129.
- Static pseudowire where each label is provisioned statically.
- Graceful Routing Engine switchover.
- Nonstop active routing.
- Multihoming.
- Partial connectivity verification (originating from an S-PE) in OAM.

RELATED DOCUMENTATION

Example: Configuring a Multisegment Pseudowire

IN THIS SECTION

- [Requirements | 1253](#)
- [Overview | 1254](#)
- [Configuration | 1260](#)
- [Verification | 1286](#)
- [Troubleshooting | 1297](#)

This example shows how to configure a dynamic multisegment pseudowire (MS-PW), where the stitching provider edge (S-PE) devices are automatically and dynamically discovered by BGP, and pseudowires are signaled by LDP using FEC 129. This arrangement requires minimum provisioning on the S-PEs, thereby reducing the configuration burden that is associated with statically configured Layer 2 circuits while still using LDP as the underlying signaling protocol.

Requirements

This example uses the following hardware and software components:

- Six routers that can be a combination of M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, T Series Core Routers, or PTX Series Packet Transport Routers.
 - Two remote PE devices configured as terminating PEs (T-PEs).
 - Two S-PEs configured as:
 - Route reflectors, in the case of interarea configuration.
 - AS boundary routers or route reflectors, in the case of inter-AS configuration.
- Junos OS Release 13.3 or later running on all the devices.

Before you begin:

1. Configure the device interfaces.
2. Configure OSPF or any other IGP protocol.
3. Configure BGP.
4. Configure LDP.
5. Configure MPLS.

Overview

Starting with Junos OS Release 13.3, you can configure an MS-PW using FEC 129 with LDP signaling and BGP autodiscovery in an MPLS packet-switched network (PSN). The MS-PW feature also provides operation, administration, and management (OAM) capabilities, such as ping, traceroute, and BFD, from the T-PE devices.

To enable autodiscovery of S-PEs in an MS-PW, include the **auto-discovery-mspw** statement at the **[edit protocols bgp group group-name family l2vpn]** hierarchy level.

```
family l2vpn {
  auto-discovery-mspw;
}
```

The automatic selection of S-PE and dynamic setting up of an MS-PW rely heavily on BGP. BGP network layer reachability information (NLRI) constructed for the FEC 129 pseudowire to autodiscover the S-PE is called an MS-PW NLRI [draft-ietf-pwe3-dynamic-ms-pw-15.txt]. The MS-PW NLRI is essentially a prefix consisting of a route distinguisher (RD) and FEC 129 source attachment identifier (SAII). It is referred to as a BGP autodiscovery (BGP-AD) route and is encoded as **RD:SAII**.

Only T-PEs that are provisioned with type 2 AIs initiate their own MS-PW NLRI respectively. Since a type 2 AI is globally unique, an MS-PW NLRI is used to identify a PE device to which the type 2 AI is provisioned. The difference between a type 1 AI and a type 2 AI requires that a new address family indicator (AFI) and subsequent address family identifier (SAFI) be defined in BGP to support an MS-PW. The proposed AFI and SAFI value pair used to identify the MS-PW NLRI is 25 and 6, respectively (pending IANA allocation).

The AFI and SAFI values support autodiscovery of S-PEs and should be configured on both T-PEs that originate the routes, and the S-PEs that participate in the signaling.

[Figure 102 on page 1255](#) illustrates an inter-area MS-PW setup between two remote PE routers—T-PE1 and T-PE2. The Provider (P) routers are P1 and P2, and the S-PE routers are S-PE1 and S-PE2. The MS-PW is established between T-PE1 and T-PE2, and all the devices belong to the same AS—AS 100. Since S-PE1 and S-PE2 belong to the same AS, they act as route reflectors and are also known as RR 1 and RR 2, respectively.

Figure 103 on page 1255 illustrates an inter-AS MS-PW setup. The MS-PW is established between T-PE1 and T-PE2, where T-PE1, P1, and S-PE1 belong to AS 1, and S-PE2, P2, and T-PE2 belong to AS 2. Since S-PE1 and S-PE2 belong to different ASs, they are configured as ASBR routers and are also known as ASBR 1 and ASBR 2, respectively.

Figure 102: Interarea Multisegment Pseudowire

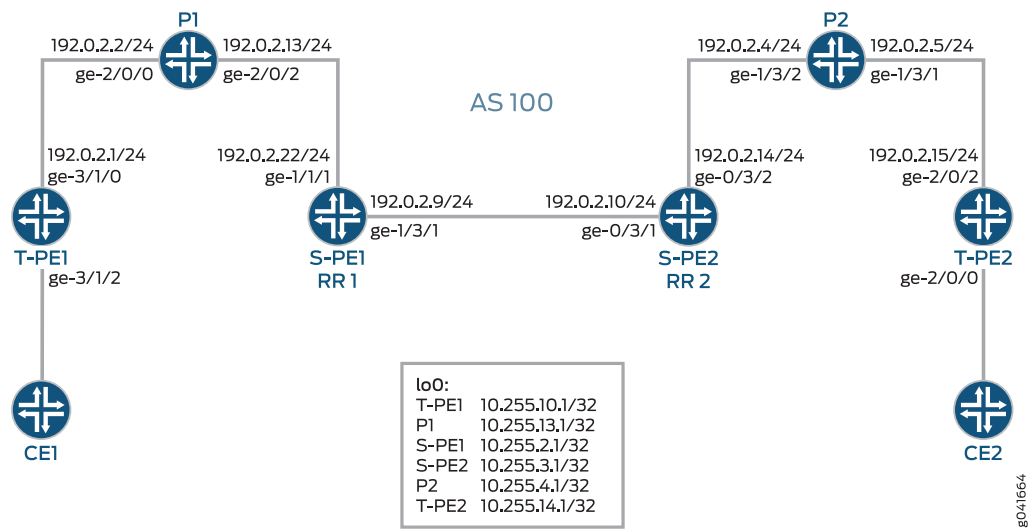
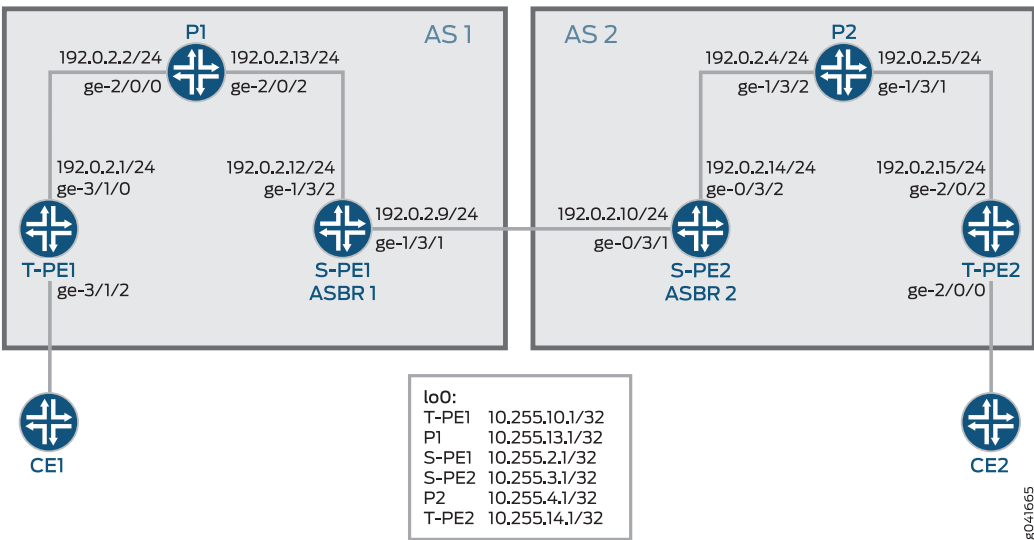


Figure 103: Inter-AS Multisegment Pseudowire



The following sections provide information about how an MS-PW is established in an interarea and inter-AS scenario.

Minimum Configuration Requirements on S-PE

In order to dynamically discover both ends of an SS-PW and set up a T-LDP session dynamically, the following is required:

- For interarea MS-PW, each S-PE plays both an ABR and BGP route reflector role.

In the interarea case, as seen in [Figure 102 on page 1255](#), the S-PE plays a BGP route reflector role and reflects the BGP-AD route to its client. A BGP-AD route advertised by one T-PE eventually reaches its remote T-PE. Because of the next-hop-self set by each S-PE, the S-PE or T-PE that receives a BGP-AD route can always discover the S-PE that advertises the BGP-AD in its local AS or local area through the BGP next hop.

- For inter-AS MS-PW, each S-PE plays either an ASBR or a BGP route reflector role.

In an MS-PW, the two T-PEs initiate a BGP-AD route respectively. When the S-PE receives the BGP-AD route through either the IBGP session with the T-PE or through a regular BGP-RR, it sets the next-hop-self before re-advertising the BGP-AD route to one or more of its EBGPs peers in the inter-AS case, as seen in [Figure 103 on page 1255](#).

- Each S-PE must set next-hop-self when re-advertising or reflecting a BGP-AD route for the MS-PW.

Active and Passive Role of T-PE

To ensure that the same set of S-PEs are being used for a MS-PW in both directions, the two T-PEs play different roles in terms of FEC 129 signaling. This is to avoid different paths being chosen by T-PE1 and T-PE2 when each S-PE is dynamically selected for an MS-PW.

When an MS-PW is signaled using FEC 129, each T-PE might independently start signaling the MS-PW. The signaling procedure can result in an attempt to set up each direction of the MS-PW through different S-PEs.

To avoid this situation, one of the T-PEs must start the pseudowire signaling (active role), while the other waits to receive the LDP label mapping before sending the respective pseudowire LDP label mapping message (passive role). When the MS-PW path is dynamically placed, the active T-PE (the Source T-PE) and the passive T-PE (the Target T-PE) must be identified before signaling is initiated for a given MS-PW. The determination of which T-PE assumes the active role is done based on the SAll value, where the T-PE that has a larger SAll value plays the active role.

In this example, the SAll values of T-PE1 and T-PE 2 are **800:800:800** and **700:700:700**, respectively. Since T-PE1 has a higher SAll value, it assumes the active role and T-PE2 assumes the passive role.

Directions for Establishing an MS-PW

The directions used by the S-PE for setting up the MS-PW are:

- Forwarding direction—From an active T-PE to a passive T-PE.

In this direction, the S-PEs perform a BGP-AD route lookup to determine the next-hop S-PE to send the label mapping message.

- Reverse direction—From a passive T-PE to an active T-PE.

In this direction, the S-PEs do not perform a BGP-AD route lookup, because the label mapping messages are received from the T-PEs, and the stitching routes are installed in the S-PEs.

In this example, the MS-PW is established in the forwarding direction from T-PE1 to T-PE2. When the MS-PW is placed from T-PE2 to T-PE1, the MS-PW is established in the reverse direction.

Autodiscovery and Dynamic Selection of S-PE

A new AFI and SAFI value is defined in BGP to support the MS-PWs based on type 2 All. This new address family supports autodiscovery of S-PEs. This address family must be configured on both the TPEs and SPEs.

It is the responsibility of the Layer 2 VPN component to dynamically select the next S-PE to use along the MS-PW in the forwarding direction.

- In the forwarding direction, the selection of the next S-PE is based on the BGP-AD route advertised by the BGP and pseudowire FEC information sent by the LDP. The BGP-AD route is initiated by the passive T-PE (T-PE2) in the reverse direction while the pseudowire FEC information is sent by LDP from the active T-PE (T-PE1) in the forwarding direction.
- In the reverse direction, the next S-PE (S-PE2) or the active T-PE (T-PE1) is obtained by looking up the S-PE (S-PE1) that it used to set up the pseudowire in the forwarding direction.

Provisioning a T-PE

To support FEC 129 type 2 All, the T-PE needs to configure its remote T-PE's IP address, a global ID, and an attachment circuit ID. Explicit paths where a set of S-PEs to use is explicitly specified on a T-PE is not supported. This eliminates the need to provision each S-PE with a type 2 All.

Stitching an MS-PW

An S-PE performs the following MPLS label operations before forwarding the received label mapping message to the next S-PE:

1. Pops the MPLS tunnel label.
2. Pops the VC label.
3. Pushes a new VC label.
4. Pushes an MPLS tunnel label used for the next segment.

Establishing an MS-PW

After completing the necessary configuration, an MS-PW is established in the following manner:

1. The SAll values are exchanged between T-PE1 and T-PE2 using BGP.
T-PE1 assumes the active T-PE role, because it is configured with a higher SAll value. T-PE2 becomes the passive T-PE.
2. T-PE1 receives the BGP-AD route originated by T-PE2. It compares the All values obtained from T-PE2 in the received BGP-AD route against the All values provisioned locally.
3. If the All values match, T-PE1 performs a BGP-AD route lookup to elect the first S-PE (S-PE1).
4. T-PE1 sends an LDP label mapping message to S-PE1.
5. Using the BGP-AD route originated from T-PE2, and the LDP label mapping message received from T-PE1, S-PE1 selects the next S-PE (S-PE2) in the forwarding direction.
To do this, S-PE1 compares SAll obtained from the BGP-AD route against the TAI from the LDP label mapping message.
6. If the All values match, S-PE1 finds S-PE2 through the BGP next hop associated with the BGP-AD route.
7. The process of selecting S-PE goes on until the last S-PE establishes a T-LDP session with T-PE2. When T-PE2 receives the LDP label mapping message from the last S-PE (S-PE2), it initiates its own label mapping message and sends it back to S-PE2.
8. When all the label mapping messages are received on S-PE1 and S-PE2, the S-PEs install the stitching routes. Thus, when the MS-PW is established in the reverse direction, the S-PEs need not perform BGP-AD route lookup to determine its next hop as it did in the forwarding direction.

OAM Support for an MS-PW

After the MS-PW is established, the following OAM capabilities can be executed from the T-PE devices:

- Ping
 - End-to-End Connectivity Verification Between T-PEs

If T-PE1, S-PEs, and T-PE2 support Control Word (CW), the pseudowire control plane automatically negotiates the use of the CW. Virtual Circuit Connectivity Verification (VCCV) Control Channel (CC) Type 3 will function correctly whether or not the CW is enabled on the pseudowire. However, VCCV Type 1, which is used for end-to-end verification only, is only supported if the CW is enabled.

The following is a sample:


```
user@T-PE1> ping mpls l2vpn fec129 instance instance-name local-id SAll of T-PE1 remote-pe-address  
address of T-PE2 remote-id TAll of T-PE2
```

or

```
user@T-PE1> ping mpls l2vpn fec129 interface CE1-facing interface
```

- Partial Connectivity Verification from T-PE to Any S-PE

To trace part of an MS-PW, the TTL of the pseudowire label can be used to force the VCCV message to pop out at an intermediate node. When the TTL expires, the S-PE can determine that the packet is a VCCV packet either by checking the CW or by checking for a valid IP header with UDP destination port 3502 (if the CW is not in use). The packet should then be diverted to VCCV processing.

If T-PE1 sends a VCCV message with the TTL of the pseudowire label equal to 1, the TTL expires at the S-PE. T-PE1 can thus verify the first segment of the pseudowire.

The VCCV packet is built according to RFC 4379. All the information necessary to build the VCCV LSP ping packet is collected by inspecting the S-PE TLVs. This use of the TTL is subject to the caution expressed in RFC 5085. If a penultimate LSR between S-PEs or between an S-PE and a T-PE manipulates the pseudowire label TTL, the VCCV message might not emerge from the MS-PW at the correct S-PE.

The following is a sample:

```
user@T-PE1> ping mpls l2vpn fec129 interface CE1-facing interface bottom-label-ttl segment
```

The **bottom-label-ttl** value is 1 for S-PE1 and 2 for S-PE2.

The **bottom-label-ttl** statement sets the correct VC label TTL, so the packets are popped to the correct SS-PW for VCCV processing.

NOTE: Junos OS supports VCCV Type 1 and Type 3 for the MS-PW OAM capability. VCCV Type 2 is not supported.

- Traceroute

Traceroute tests each S-PE along the path of the MS-PW in a single operation similar to LSP trace. This operation is able to determine the actual data path of the MS-PW, and is used for dynamically signaled MS-PWs.

```
user@T-PE1> traceroute mpls l2vpn fec129 interface CE1-facing interface
```

- Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. In addition to fast forwarding path failure detection, BFD provides a consistent failure detection method

for network administrators. The router or switch can be configured to log a system log (syslog) message when BFD goes down.

```
user@T-PE1> show bfd session extensive
```

Configuration

IN THIS SECTION

- [Configuring an Interarea MS-PW | 1260](#)
- [Configuring an Inter-AS MS-PW | 1273](#)

Configuring an Interarea MS-PW

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.

T-PE1

```
set interfaces ge-3/1/0 unit 0 family inet address 192.0.2.1/24
set interfaces ge-3/1/0 unit 0 family mpls
set interfaces ge-3/1/2 encapsulation ethernet-ccc
set interfaces ge-3/1/2 unit 0
set interfaces lo0 unit 0 family inet address 10.255.10.1/32 primary
set routing-options autonomous-system 100
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw
set protocols bgp group mspw type internal
set protocols bgp group mspw local-address 10.255.10.1
set protocols bgp group mspw neighbor 10.255.2.1
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
```



```

set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set routing-instances ms-pw instance-type l2vpn
set routing-instances ms-pw interface ge-3/1/2.0
set routing-instances ms-pw route-distinguisher 10.10.10.10:15
set routing-instances ms-pw l2vpn-id l2vpn-id:100:15
set routing-instances ms-pw vrf-target target:100:115
set routing-instances ms-pw protocols l2vpn site CE1 source-attachment-identifier 800:800:800
set routing-instances ms-pw protocols l2vpn site CE1 interface ge-3/1/2.0 target-attachment-identifier
    700:700:700
set routing-instances ms-pw protocols l2vpn pseudowire-status-tlv
set routing-instances ms-pw protocols l2vpn oam bfd-liveness-detection minimum-interval 300

```

P1

```

set interfaces ge-2/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.13/24
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.13.1/32 primary
set routing-options autonomous-system 100
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0

```

S-PE1 (RR 1)

```

set interfaces ge-1/3/1 unit 0 family inet address 192.0.2.9/24
set interfaces ge-1/3/1 unit 0 family mpls
set interfaces ge-1/3/2 unit 0 family inet address 192.0.2.22/24
set interfaces ge-1/3/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.2.1/32 primary

```



```

set routing-options autonomous-system 100
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw
set protocols bgp group mspw type internal
set protocols bgp group mspw local-address 10.255.2.1
set protocols bgp group mspw export next-hop-self
set protocols bgp group mspw cluster 203.0.113.0
set protocols bgp group mspw neighbor 10.255.10.1
set protocols bgp group mspw neighbor 10.255.3.1
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-inet0 from protocol bgp
set policy-options policy-statement send-inet0 then accept

```

S-PE2 (RR 2)

```

set interfaces ge-0/3/1 unit 0 family inet address 192.0.2.10/24
set interfaces ge-0/3/1 unit 0 family mpls
set interfaces ge-0/3/2 unit 0 family inet address 192.0.2.14/24
set interfaces ge-0/3/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.3.1/32 primary
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw
set protocols bgp group mspw type internal
set protocols bgp group mspw local-address 10.255.3.1
set protocols bgp group mspw export next-hop-self
set protocols bgp group mspw cluster 198.51.100.0
set protocols bgp group mspw neighbor 10.255.2.1
set protocols bgp group mspw neighbor 10.255.14.1
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.3.1
set protocols bgp group int neighbor 10.255.2.1
set protocols ospf area 0.0.0.0 interface all

```



```

set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-inet0 from protocol bgp
set policy-options policy-statement send-inet0 then accept

```

P2

```

set interfaces ge-1/3/1 unit 0 family inet address 192.0.2.5/24
set interfaces ge-1/3/1 unit 0 family mpls
set interfaces ge-1/3/2 unit 0 family inet address 192.0.2.4/24
set interfaces ge-1/3/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.4.1/32 primary
set routing-options autonomous-system 100
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0

```

T-PE2

```

set interfaces ge-2/0/0 encapsulation ethernet-ccc
set interfaces ge-2/0/0 unit 0
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.15/24
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.14.1/32 primary
set routing-options autonomous-system 100
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw

```



```

set protocols bgp group mspw type internal
set protocols bgp group mspw local-address 10.255.14.1
set protocols bgp group mspw neighbor 10.255.3.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set routing-instances ms-pw instance-type l2vpn
set routing-instances ms-pw interface ge-2/0/0.0
set routing-instances ms-pw route-distinguisher 10.10.10.10:15
set routing-instances ms-pw l2vpn-id l2vpn-id:100:15
set routing-instances ms-pw vrf-target target:100:115
set routing-instances ms-pw protocols l2vpn site CE2 source-attachment-identifier 700:700:700
set routing-instances ms-pw protocols l2vpn site CE2 interface ge-2/0/0.0 target-attachment-identifier
  800:800:800
set routing-instances ms-pw protocols l2vpn pseudowire-status-tlv
set routing-instances ms-pw protocols l2vpn oam bfd-liveness-detection minimum-interval 300

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure T-PE1 in the interarea scenario:

NOTE: Repeat this procedure for the T-PE2 device in the MPLS domain, after modifying the appropriate interface names, addresses, and other parameters.

1. Configure the T-PE1 interfaces.

```

[edit interfaces]
user@T-PE1# set ge-3/1/0 unit 0 family inet address 192.0.2.1/24
user@T-PE1# set ge-3/1/0 unit 0 family mpls
user@T-PE1# set ge-3/1/2 encapsulation ethernet-ccc
user@T-PE1# set ge-3/1/2 unit 0
user@T-PE1# set lo0 unit 0 family inet address 10.255.10.1/32 primary

```

2. Set the autonomous system number.


```
[edit routing-options]
user@T-PE1# set autonomous-system 100
```

3. Enable MPLS on all the interfaces of T-PE1, excluding the management interface.

```
[edit protocols]
user@T-PE1# set mpls interface all
user@T-PE1# set mpls interface fxp0.0 disable
```

4. Enable autodiscovery of intermediate S-PEs that make up the MS-PW using BGP.

```
[edit protocols]
user@T-PE1# set bgp family l2vpn auto-discovery-mspw
```

5. Configure the BGP group for T-PE1.

```
[edit protocols]
user@T-PE1# set bgp group mspw type internal
```

6. Assign local and neighbor addresses to the mspw group for T-PE1 to peer with S-PE1.

```
[edit protocols]
user@T-PE1# set bgp group mspw local-address 10.255.10.1
user@T-PE1# set bgp group mspw neighbor 10.255.2.1
```

7. Configure OSPF on all the interfaces of T-PE1, excluding the management interface.

```
[edit protocols]
user@T-PE1# set ospf area 0.0.0.0 interface lo0.0
user@T-PE1# set ospf area 0.0.0.0 interface all
user@T-PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
```

8. Configure LDP on all the interfaces of T-PE1, excluding the management interface.

```
[edit protocols]
user@T-PE1# set ldp interface all
user@T-PE1# set ldp interface fxp0.0 disable
user@T-PE1# set ldp interface lo0.0
```


9. Configure the Layer 2 VPN routing instance on T-PE1.

```
[edit routing-instances]
user@T-PE1# set ms-pw instance-type l2vpn
```

10. Assign the interface name for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw interface ge-3/1/2.0
```

11. Configure the route distinguisher for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw route-distinguisher 10.10.10.15
```

12. Configure the Layer 2 VPN ID community for FEC 129 MS-PW.

```
[edit routing-instances]
user@T-PE1# set ms-pw l2vpn-id l2vpn-id:100:15
```

13. Configure a VPN routing and forwarding (VRF) target for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw vrf-target target:100:115
```

14. Configure the source attachment identifier (SAI) value using Layer 2 VPN as the routing protocol for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn site CE1 source-attachment-identifier 800:800:800
```

15. Assign the interface name that connects the CE1 site to the VPN, and configure the target attachment identifier (TAI) value using Layer 2 VPN as the routing protocol for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn site CE1 interface ge-3/1/2.0 target-attachment-identifier
700:700:700
```


16. (Optional) Configure T-PE1 to send MS-PW status TLVs.

```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn pseudowire-status-tlv
```

17. (Optional) Configure OAM capabilities for the VPN.

```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn oam bfd-liveness-detection minimum-interval 300
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure S-PE1 (RR 1) in the interarea scenario:

NOTE: Repeat this procedure for the S-PE2 (RR 2) device in the MPLS domain, after modifying the appropriate interface names, addresses, and other parameters.

1. Configure the S-PE1 interfaces.

```
[edit interfaces]
user@S-PE1# set ge-1/3/1 unit 0 family inet address 192.0.2.9/24
user@S-PE1# set ge-1/3/1 unit 0 family mpls
user@S-PE1# set ge-1/3/2 unit 0 family inet address 192.0.2.22/24
user@S-PE1# set ge-1/3/2 unit 0 family mpls
user@S-PE1# set lo0 unit 0 family inet address 10.255.2.1/32 primary
```

2. Set the autonomous system number.

```
[edit routing-options]
user@S-PE1# set autonomous-system 100
```

3. Enable MPLS on all the interfaces of T-PE1, excluding the management interface.

```
[edit protocols]
user@S-PE1# set mpls interface all
user@S-PE1# set mpls interface fxp0.0 disable
```


4. Enable autodiscovery of S-PE using BGP.

```
[edit protocols]
user@S-PE1# set bgp family l2vpn auto-discovery-mspw
```

5. Configure the BGP group for S-PE1.

```
[edit protocols]
user@S-PE1# set bgp group mspw type internal
```

6. Configure S-PE1 to act as a route reflector.

```
[edit protocols]
user@S-PE1# set bgp group mspw export next-hop-self
user@S-PE1# set bgp group mspw cluster 203.0.113.0
```

7. Assign local and neighbor addresses to the mspw group for S-PE1 to peer with T-PE1 and S-PE2.

```
[edit protocols]
user@S-PE1# set bgp group mspw local-address 10.255.2.1
user@S-PE1# set bgp group mspw neighbor 10.255.10.1 (to T-PE1)
user@S-PE1# set bgp group mspw neighbor 10.255.3.1 (to S-PE2)
```

8. Configure OSPF on all the interfaces of S-PE1, excluding the management interface.

```
[edit protocols]
user@S-PE1# set ospf area 0.0.0.0 interface all
user@S-PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@S-PE1# set ospf area 0.0.0.0 interface lo0.0
```

9. Configure LDP on all the interfaces of S-PE1, excluding the management interface.

```
[edit protocols]
user@S-PE1# set ldp interface all
user@S-PE1# set ldp interface fxp0.0 disable
user@S-PE1# set ldp interface lo0.0
```

10. Define the policy for enabling next-hop-self and accepting BGP traffic on S-PE1.


```
[edit policy-options]
user@S-PE1# set policy-statement next-hop-self then next-hop self
user@S-PE1# set policy-statement send-inet0 from protocol bgp
user@S-PE1# set policy-statement send-inet0 then accept
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-instances**, **show routing-options**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

T-PE1

```
user@T-PE1# show interfaces
ge-3/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
    family mpls;
  }
}
ge-3/1/2 {
  encapsulation ethernet-ccc;
  unit 0;
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.10.1/32 {
        primary;
      }
    }
  }
}
```

```
user@T-PE1# show routing-options
autonomous-system 100;
```

```
user@T-PE1# show protocols
```



```

mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    family l2vpn {
        auto-discovery-mspw;
    }
    group mspw {
        type internal;
        local-address 10.255.10.1;
        neighbor 10.255.2.1;
    }
}
ospf {
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}

```

```

user@T-PE1# show routing-instances
ms-pw {
    instance-type l2vpn;
    interface ge-3/1/2.0;
    route-distinguisher 10.10.10.10:15;
    l2vpn-id l2vpn-id:100:15;
    vrf-target target:100:115;
    protocols {
        l2vpn {
            site CE1 {
                source-attachment-identifier 800:800:800;
            }
        }
    }
}

```



```
user@S-PE1# show interfaces
ge-1/3/1 {
  unit 0 {
    family inet {
      address 192.0.2.9/24;
    }
    family mpls;
  }
}
ge-1/3/2 {
  unit 0 {
    family inet {
      address 192.0.2.22/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.2.1/32 {
        primary;
      }
    }
  }
}
```



```
user@S-PE1# show routing-options
autonomous-system 100;
```

```
user@S-PE1# show protocols
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  family l2vpn {
    auto-discovery-mspw;
  }
  group mspw {
    type internal;
    local-address 10.255.2.1;
    export next-hop-self;
    cluster 203.0.113.0;
    neighbor 10.255.10.1;
    neighbor 10.255.3.1;
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0;
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
ldp {
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0;
}
```

```
user@S-PE1# show policy-options
policy-statement next-hop-self {
  then {
```



```

        next-hop self;
    }
}
policy-statement send-inet0 {
    from protocol bgp;
    then accept;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring an Inter-AS MS-PW

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.

T-PE1

```

set interfaces ge-3/1/0 unit 0 family inet address 192.0.2.1/24
set interfaces ge-3/1/0 unit 0 family mpls
set interfaces ge-3/1/2 encapsulation ethernet-ccc
set interfaces ge-3/1/2 unit 0
set interfaces lo0 unit 0 family inet address 10.255.10.1/32 primary
set routing-options autonomous-system 1
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw
set protocols bgp group mspw type internal
set protocols bgp group mspw local-address 10.255.10.1
set protocols bgp group mspw neighbor 10.255.2.1
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set routing-instances ms-pw instance-type l2vpn
set routing-instances ms-pw interface ge-3/1/2.0
set routing-instances ms-pw route-distinguisher 10.10.10.10:15
set routing-instances ms-pw l2vpn-id l2vpn-id:100:15
set routing-instances ms-pw vrf-target target:100:115

```



```

set routing-instances ms-pw protocols l2vpn site CE1 source-attachment-identifier 800:800:800
set routing-instances ms-pw protocols l2vpn site CE1 interface ge-3/1/2.0 target-attachment-identifier
  700:700:700
set routing-instances ms-pw protocols l2vpn pseudowire-status-tlv
set routing-instances ms-pw protocols l2vpn oam bfd-liveness-detection minimum-interval 300

```

P1

```

set interfaces ge-2/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-2/0/0 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.13/24
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.13.1/32 primary
set routing-options autonomous-system 1
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0

```

S-PE1 (ASBR 1)

```

set interfaces ge-1/3/1 unit 0 family inet address 192.0.2.9/24
set interfaces ge-1/3/1 unit 0 family mpls
set interfaces ge-1/3/2 unit 0 family inet address 192.0.2.22/24
set interfaces ge-1/3/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.2.1/32 primary
set routing-options autonomous-system 1
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw
set protocols bgp group to_T-PE1 type internal
set protocols bgp group to_T-PE1 local-address 10.255.2.1
set protocols bgp group to_T-PE1 export next-hop-self

```



```

set protocols bgp group to_T-PE1 neighbor 10.255.10.1
set protocols bgp group to_S-PE2 type external
set protocols bgp group to_S-PE2 local-address 10.255.2.1
set protocols bgp group to_S-PE2 peer-as 2
set protocols bgp group to_S-PE2 neighbor 10.255.3.1 multihop ttl 1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement next-hop-self then next-hop self

```

S-PE2 (ASBR 2)

```

set interfaces ge-0/3/1 unit 0 family inet address 192.0.2.10/24
set interfaces ge-0/3/1 unit 0 family mpls
set interfaces ge-0/3/2 unit 0 family inet address 192.0.2.14/24
set interfaces ge-0/3/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.3.1/32 primary
set routing-options autonomous-system 2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw
set protocols bgp group to_T-PE2 type internal
set protocols bgp group to_T-PE2 local-address 10.255.3.1
set protocols bgp group to_T-PE2 export next-hop-self
set protocols bgp group to_T-PE2 neighbor 10.255.14.1
set protocols bgp group to_S-PE1 type external
set protocols bgp group to_S-PE1 local-address 10.255.3.1
set protocols bgp group to_S-PE1 peer-as 1
set protocols bgp group to_S-PE1 neighbor 10.255.2.1 multihop ttl 1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set policy-options policy-statement next-hop-self then next-hop self

```


P2

```

set interfaces ge-1/3/1 unit 0 family inet address 192.0.2.5/24
set interfaces ge-1/3/1 unit 0 family mpls
set interfaces ge-1/3/2 unit 0 family inet address 192.0.2.4/24
set interfaces ge-1/3/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.4.1/32 primary
set routing-options autonomous-system 2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0

```

T-PE2

```

set interfaces ge-2/0/0 encapsulation ethernet-ccc
set interfaces ge-2/0/0 unit 0
set interfaces ge-2/0/2 unit 0 family inet address 192.0.2.15/24
set interfaces ge-2/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.14.1/32 primary
set routing-options autonomous-system 2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp family l2vpn auto-discovery-mspw
set protocols bgp group mspw type internal
set protocols bgp group mspw local-address 10.255.14.1
set protocols bgp group mspw neighbor 10.255.3.1
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set protocols ldp interface lo0.0
set routing-instances ms-pw instance-type l2vpn
set routing-instances ms-pw interface ge-2/0/0.0
set routing-instances ms-pw route-distinguisher 10.10.10.10:15

```



```

set routing-instances ms-pw l2vpn-id l2vpn-id:100:15
set routing-instances ms-pw vrf-target target:100:115
set routing-instances ms-pw protocols l2vpn site CE2 source-attachment-identifier 700:700:700
set routing-instances ms-pw protocols l2vpn site CE2 interface ge-2/0/0.0 target-attachment-identifier
  800:800:800
set routing-instances ms-pw protocols l2vpn pseudowire-status-tlv
set routing-instances ms-pw protocols l2vpn oam bfd-liveness-detection minimum-interval 300

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the T-PE1 router in the inter-AS scenario:

NOTE: Repeat this procedure for the T-PE2 device in the MPLS domain, after modifying the appropriate interface names, addresses, and other parameters.

1. Configure the T-PE1 interfaces.

```

[edit interfaces]
user@T-PE1# set ge-3/1/0 unit 0 family inet address 192.0.2.1/24
user@T-PE1# set ge-3/1/0 unit 0 family mpls
user@T-PE1# set ge-3/1/2 encapsulation ethernet-ccc
user@T-PE1# set ge-3/1/2 unit 0
user@T-PE1# set lo0 unit 0 family inet address 10.255.10.1/32 primary

```

2. Set the autonomous system number.

```

[edit routing-options]
user@T-PE1# set autonomous-system 1

```

3. Enable MPLS on all the interfaces of T-PE1, excluding the management interface.

```

[edit protocols]
user@T-PE1# set mpls interface all
user@T-PE1# set mpls interface fxp0.0 disable

```


4. Enable autodiscovery of intermediate S-PEs that make up the MS-PW using BGP.

```
[edit protocols]
user@T-PE1# set bgp family l2vpn auto-discovery-mspw
```

5. Configure the BGP group for T-PE1.

```
[edit protocols]
user@T-PE1# set bgp group mspw type internal
```

6. Assign local and neighbor addresses to the mspw group for T-PE1 to peer with S-PE1.

```
[edit protocols]
user@T-PE1# set bgp group mspw local-address 10.255.10.1
user@T-PE1# set bgp group mspw neighbor 10.255.2.1
```

7. Configure OSPF on all the interfaces of T-PE1, excluding the management interface.

```
[edit protocols]
user@T-PE1# set ospf area 0.0.0.0 interface lo0.0
user@T-PE1# set ospf area 0.0.0.0 interface all
user@T-PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
```

8. Configure LDP on all the interfaces of T-PE1, excluding the management interface.

```
[edit protocols]
user@T-PE1# set ldp interface all
user@T-PE1# set ldp interface fxp0.0 disable
user@T-PE1# set ldp interface lo0.0
```

9. Configure the Layer 2 VPN routing instance on T-PE1.

```
[edit routing-instances]
user@T-PE1# set ms-pw instance-type l2vpn
```

10. Assign the interface name for the mspw routing instance.

```
[edit routing-instances]
```



```
user@T-PE1# set ms-pw interface ge-3/1/2.0
```

11. Configure the route distinguisher for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw route-distinguisher 10.10.10.15
```

12. Configure the Layer 2 VPN ID community for FEC 129 MS-PW.

```
[edit routing-instances]
user@T-PE1# set ms-pw l2vpn-id l2vpn-id:100:15
```

13. Configure a VPN routing and forwarding (VRF) target for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw vrf-target target:100:115
```

14. Configure the source attachment identifier (SAI) value using Layer 2 VPN as the routing protocol for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn site CE1 source-attachment-identifier 800:800:800
```

15. Assign the interface name that connects the CE1 site to the VPN, and configure the target attachment identifier (TAI) value using Layer 2 VPN as the routing protocol for the mspw routing instance.

```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn site CE1 interface ge-3/1/2.0 target-attachment-identifier
700:700:700
```

16. (Optional) Configure T-PE1 to send MS-PW status TLVs.

```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn pseudowire-status-tlv
```

17. (Optional) Configure OAM capabilities for the VPN.


```
[edit routing-instances]
user@T-PE1# set ms-pw protocols l2vpn oam bfd-liveness-detection minimum-interval 300
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure S-PE1 (ASBR 1) in the inter-AS scenario:

NOTE: Repeat this procedure for the S-PE2 (ASBR 2) device in the MPLS domain, after modifying the appropriate interface names, addresses, and other parameters.

1. Configure S-PE1 (ASBR 1) interfaces.

```
[edit interfaces]
user@S-PE1# set ge-1/3/1 unit 0 family inet address 192.0.2.9/24
user@S-PE1# set ge-1/3/1 unit 0 family mpls
user@S-PE1# set ge-1/3/2 unit 0 family inet address 192.0.2.22/24
user@S-PE1# set ge-1/3/2 unit 0 family mpls
user@S-PE1# set lo0 unit 0 family inet address 10.255.2.1/32 primary
```

2. Set the autonomous system number.

```
[edit routing-options]
user@S-PE1# set autonomous-system 1
```

3. Enable MPLS on all the interfaces of S-PE1 (ASBR 1), excluding the management interface.

```
[edit protocols]
user@S-PE1# set mpls interface all
user@S-PE1# set mpls interface fxp0.0 disable
```

4. Enable autodiscovery of S-PE using BGP.

```
[edit protocols]
user@S-PE1# set bgp family l2vpn auto-discovery-mspw
```


5. Configure the IBGP group for S-PE1 (ASBR 1) to peer with T-PE1.

```
[edit protocols]
user@S-PE1# set bgp group to_T-PE1 type internal
```

6. Configure the IBGP group parameters.

```
[edit protocols]
user@S-PE1# set bgp group to_T-PE1 local-address 10.255.2.1
user@S-PE1# set bgp group to_T-PE1 export next-hop-self
user@S-PE1# set bgp group to_T-PE1 neighbor 10.255.10.1
```

7. Configure the EBGP group for S-PE1 (ASBR 1) to peer with S-PE2 (ASBR 2).

```
[edit protocols]
user@S-PE1# set bgp group to_S-PE2 type external
```

8. Configure the EBGP group parameters.

```
[edit protocols]
user@S-PE1# set bgp group to_S-PE2 local-address 10.255.2.1
user@S-PE1# set bgp group to_S-PE2 peer-as 2
user@S-PE1# set bgp group to_S-PE2 neighbor 10.255.3.1 multihop ttl 1
```

9. Configure OSPF on all the interfaces of S-PE1 (ASBR 1), excluding the management interface.

```
[edit protocols]
user@S-PE1# set ospf area 0.0.0.0 interface all
user@S-PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@S-PE1# set ospf area 0.0.0.0 interface lo0.0 passive
```

10. Configure LDP on all the interfaces of S-PE1 (ASBR 1), excluding the management interface.

```
[edit protocols]
user@S-PE1# set ldp interface all
user@S-PE1# set ldp interface fxp0.0 disable
user@S-PE1# set ldp interface lo0.0
```

11. Define the policy for enabling next-hop-self on S-PE1 (ASBR 1).


```
[edit policy-options]
user@S-PE1# set policy-statement next-hop-self then next-hop self
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show routing-instances**, **show routing-options**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

T-PE1

```
user@T-PE1# show interfaces
ge-3/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
    family mpls;
  }
}
ge-3/1/2 {
  encapsulation ethernet-ccc;
  unit 0;
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.10.1/32 {
        primary;
      }
    }
  }
}
```

```
user@T-PE1# show routing-options
autonomous-system 1;
```

```
user@T-PE1# show protocols
mpls {
  interface all;
```



```

interface fxp0.0 {
    disable;
}
}
bgp {
    family l2vpn {
        auto-discovery-mspw;
    }
    group mspw {
        type internal;
        local-address 10.255.10.1;
        neighbor 10.255.2.1;
    }
}
ospf {
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
        interface lo0.0;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}

```

```

user@T-PE1# show routing-instances
ms-pw {
    instance-type l2vpn;
    interface ge-3/1/2.0;
    route-distinguisher 10.10.10.10:15;
    l2vpn-id l2vpn-id:100:15;
    vrf-target target:100:115;
    protocols {
        l2vpn {
            site CE1 {
                source-attachment-identifier 800:800:800;
                interface ge-3/1/2.0 {
                    target-attachment-identifier 700:700:700;
                }
            }
        }
    }
}

```



```

    }
  }
  pseudowire-status-tlv;
  oam {
    bfd-liveness-detection {
      minimum-interval 300;
    }
  }
}
}
}
```

S-PE1 (RR 1)

```
user@S-PE1# show interfaces
ge-1/3/1 {
  unit 0 {
    family inet {
      address 192.0.2.9/24;
    }
    family mpls;
  }
}
ge-1/3/2 {
  unit 0 {
    family inet {
      address 192.0.2.22/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.2.1/32 {
        primary;
      }
    }
  }
}
```



```
user@T-PE1# show routing-options
autonomous-system 1;
```

```
user@S-PE1# show protocols
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  family l2vpn {
    auto-discovery-mspw;
  }
  group to_T-PE1 {
    type internal;
    local-address 10.255.2.1;
    export next-hop-self;
    neighbor 10.255.10.1;
  }
  group to_S-PE2 {
    type external;
    local-address 10.255.2.1;
    peer-as 2;
    neighbor 10.255.3.1 {
      multihop {
        ttl 1;
      }
    }
  }
}
ospf {
  area 0.0.0.0 {
    interface lo0.0 {
      passive;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
ldp {
  interface all;
```



```

interface fxp0.0 {
  disable;
}
interface lo0.0;
}

```

```

user@T-PE1# show policy-options
policy-statement next-hop-self {
  then {
    next-hop self;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying the Routes | 1286](#)
- [Verifying the LDP Database | 1289](#)
- [Checking the MS-PW Connections on T-PE1 | 1290](#)
- [Checking the MS-PW Connections on S-PE1 | 1292](#)
- [Checking the MS-PW Connections on S-PE2 | 1293](#)
- [Checking the MS-PW Connections on T-PE2 | 1295](#)

Confirm that the configuration is working properly.

Verifying the Routes

Purpose

Verify that the expected routes are learned.

Action

From operational mode, run the **show route** command for the **bgp.l2vpn.1**, **ldp.l2vpn.1**, **mpls.0**, and **ms-pw.l2vpn.1** routing tables.

From operational mode, run the **show route table bgp.l2vpn.1** command.

user@T-PE1> **show route table *bgp.l2vpn.1***

```
bgp.l2vpn.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

10.10.10.10:15:700:0.0.2.188:700/160 AD2

```
*[BGP/170] 16:13:11, localpref 100, from 10.255.2.1
  AS path: 2 I, validation-state: unverified
  > to 203.0.113.2 via ge-3/1/0.0, Push 300016
```

From operational mode, run the **show route table *ldp.l2vpn.1*** command.

user@T-PE1> **show route table *ldp.l2vpn.1***

```
ldp.l2vpn.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

10.255.2.1:CtrlWord:5:100:15:700:0.0.2.188:700:800:0.0.3.32:800/304 PW2

```
*[LDP/9] 16:21:27
  Discard
```

From operational mode, run the **show route table *mpls.0*** command.

user@T-PE1> **show route table *mpls.0***

```
mpls.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
0          *[MPLS/0] 1w6d 00:28:26, metric 1
           Receive
1          *[MPLS/0] 1w6d 00:28:26, metric 1
           Receive
2          *[MPLS/0] 1w6d 00:28:26, metric 1
           Receive
13         *[MPLS/0] 1w6d 00:28:26, metric 1
           Receive
299920     *[LDP/9] 1w5d 01:26:08, metric 1
           > to 203.0.113.2 via ge-3/1/0.0, Pop
299920(S=0) *[LDP/9] 1w5d 01:26:08, metric 1
           > to 203.0.113.2 via ge-3/1/0.0, Pop
299936     *[LDP/9] 1w5d 01:26:08, metric 1
           > to 203.0.113.2 via ge-3/1/0.0, Swap 300016
```



```

300096      *[LDP/9] 16:22:35, metric 1
           > to 203.0.113.2 via ge-3/1/0.0, Swap 300128
300112      *[LDP/9] 16:22:35, metric 1
           > to 203.0.113.2 via ge-3/1/0.0, Swap 300144
300128      *[LDP/9] 16:22:35, metric 1
           > to 203.0.113.2 via ge-3/1/0.0, Swap 300160
300144      *[L2VPN/7] 16:22:33
           > via ge-3/1/2.0, Pop      Offset: 4
ge-3/1/2.0  *[L2VPN/7] 16:22:33, metric2 1
           > to 203.0.113.2 via ge-3/1/0.0, Push 300176, Push 300016(top)

Offset: 252

```

From operational mode, run the **show route table ms-pw.l2vpn.1** command.

user@T-PE1> **show route table ms-pw.l2vpn.1**

```

ms-pw.l2vpn.1: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.10.10:15:700:0.0.2.188:700/160 AD2
      *[BGP/170] 16:23:27, localpref 100, from 10.255.2.1
      AS path: 2 I, validation-state: unverified
      > to 203.0.113.2 via ge-3/1/0.0, Push 300016
10.10.10.10:15:800:0.0.3.32:800/160 AD2
      *[L2VPN/170] 1w5d 23:25:19, metric2 1
      Indirect
10.255.2.1:CtrlWord:5:100:15:700:0.0.2.188:700:800:0.0.3.32:800/304 PW2

      *[LDP/9] 16:23:25
      Discard
10.255.2.1:CtrlWord:5:100:15:800:0.0.3.32:800:700:0.0.2.188:700/304 PW2

      *[L2VPN/7] 16:23:27, metric2 1
      > to 203.0.113.2 via ge-3/1/0.0, Push 300016

```

Meaning

The output shows all the learned routes, including the autodiscovery (AD) routes.

The AD2 prefix format is **RD:SAIL-type2**, where:

- **RD** is the route distinguisher value.
- **SAIL-type2** is the type 2 source attachment identifier value.

The PW2 prefix format is **Neighbor_Addr:C:PWtype:l2vpn-id:SAll-type2:TAIl-type2**, where:

- **Neighbor_Addr** is the loopback address of neighboring S-PE device.
- **C** indicates if Control Word (CW) is enabled or not.
 - **C** is **CtrlWord** if CW is set.
 - **C** is **NoCtrlWord** if CW is not set.
- **PWtype** indicates the type of the pseudowire.
 - **PWtype** is **4** if it is in Ethernet tagged mode.
 - **PWtype** is **5** if it is Ethernet only.
- **l2vpn-id** is the Layer 2 VPN ID for the MS-PW routing instance.
- **SAll-type2** is the type 2 source attachment identifier value.
- **TAIl-type2** is the type 2 target attachment identifier value.

Verifying the LDP Database

Purpose

Verify the MS-PW labels received by T-PE1 from S-PE1 and sent from T-PE1 to S-PE1.

Action

From operational mode, run the **show ldp database** command.

user@T-PE1> **show ldp database**

```
Input label database, 10.255.10.1:0--10.255.2.1:0
  Label      Prefix
    3        10.255.2.1/32
  300112     10.255.3.1/32
  300128     10.255.4.1/32
  299968     10.255.10.1/32
  299904     10.255.13.1/32
  300144     10.255.14.1/32
300176    FEC129 CtrlWord ETHERNET 000a0064:0000000f 000002bc:000002bc:000002bc
00000320:00000320:00000320
```

```
Output label database, 10.255.10.1:0--10.255.2.1:0
  Label      Prefix
  299936     10.255.2.1/32
  300096     10.255.3.1/32
  300112     10.255.4.1/32
```



```

      3      10.255.10.1/32
299920      10.255.13.1/32
300128      10.255.14.1/32
300144    FEC129 CtrlWord ETHERNET 000a0064:0000000f 00000320:00000320:00000320
000002bc:000002bc:000002bc

Input label database, 10.255.10.1:0--10.255.13.1:0
  Label      Prefix
300016      10.255.2.1/32
300128      10.255.3.1/32
300144      10.255.4.1/32
300080      10.255.10.1/32
      3      10.255.13.1/32
300160      10.255.14.1/32

Output label database, 10.255.10.1:0--10.255.13.1:0
  Label      Prefix
299936      10.255.2.1/32
300096      10.255.3.1/32
300112      10.255.4.1/32
      3      10.255.10.1/32
299920      10.255.13.1/32
300128      10.255.14.1/32

```

Meaning

The labels with **FEC129** prefix are related to the MS-PW.

Checking the MS-PW Connections on T-PE1

Purpose

Make sure that all of the FEC 129 MS-PW connections come up correctly.

Action

From operational mode, run the **show l2vpn connections extensive** command.

```
user@T-PE1> show l2vpn connections extensive
```

```
Layer-2 VPN connections:
```

```
Legend for connection status (St)
```

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same

```


VC-Dn -- Virtual circuit down NP -- interface hardware not present
 CM -- control-word mismatch -> -- only outbound connection is up
 CN -- circuit not provisioned <- -- only inbound connection is up
 OR -- out of range Up -- operational
 OL -- no outgoing label Dn -- down
 LD -- local site signaled down CF -- call admission control failure
 RD -- remote site signaled down SC -- local and remote site ID collision
 LN -- local site not designated LM -- local site ID not minimum designated
 RN -- remote site not designated RM -- remote site ID not minimum designated
 XX -- unknown connection status IL -- no incoming label
 MM -- MTU mismatch MI -- Mesh-Group ID not available
 BK -- Backup connection ST -- Standby connection
 PF -- Profile parse failure PB -- Profile busy
 RS -- remote site standby SN -- Static Neighbor
 LB -- Local site not best-site RB -- Remote site not best-site
 VM -- VLAN ID mismatch

Legend for interface status

Up -- operational
 Dn -- down

Instance: ms-pw

L2vpn-id: 100:15

Number of local interfaces: 1

Number of local interfaces up: 1

ge-3/1/2.0

Local source-attachment-id: 800:0.0.3.32:800 (CE1)

Target-attachment-id	Type	St	Time last up	# Up trans
700:0.0.2.188:700	rmt	Up	Sep 18 01:10:55 2013	1

Remote PE: 10.255.2.1, Negotiated control-word: Yes (Null)

Incoming label: 300048, Outgoing label: 300016

Negotiated PW status TLV: Yes

local PW status code: 0x00000000, Neighbor PW status code: 0x00000000

Local interface: ge-3/1/2.0, Status: Up, Encapsulation: ETHERNET

Pseudowire Switching Points :

Local address	Remote address	Status
10.255.2.1	10.255.3.1	forwarding
10.255.3.1	10.255.14.1	forwarding

Connection History:

Sep 18 01:10:55 2013	status update timer	
Sep 18 01:10:55 2013	PE route changed	
Sep 18 01:10:55 2013	Out lbl Update	300016
Sep 18 01:10:55 2013	In lbl Update	300048
Sep 18 01:10:55 2013	loc intf up	ge-3/1/2.0

Check the following fields in the output to verify that MS-PW is established between the T-PE devices:

- **Target-attachment-id**—Check if the TAI value is the SAI value of T-PE2.
- **Remote PE**—Check if the T-PE2 loopback address is listed.
- **Negotiated PW status TLV**—Ensure that the value is **Yes**.
- **Pseudowire Switching Points**—Check if the switching points are listed from S-PE1 to S-PE2 and from S-PE2 to T-PE2.

Meaning

MS-PW is established between T-PE1 and T-PE2 in the forwarding direction.

Checking the MS-PW Connections on S-PE1

Purpose

Make sure that all of the FEC 129 MS-PW connections come up correctly for the mspw routing instance.

Action

From operational mode, run the **show l2vpn connections instance __MSPW__ extensive** command.

```
user@S-PE1> show l2vpn connections instance __MSPW__ extensive
```

```
Layer-2 VPN connections:
```

```
Legend for connection status (St)
```

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unknown connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy
RS -- remote site standby	SN -- Static Neighbor
LB -- Local site not best-site	RB -- Remote site not best-site
VM -- VLAN ID mismatch	


```

Legend for interface status
Up -- operational
Dn -- down

Instance: __MSPW__
  L2vpn-id: 100:15
  Local source-attachment-id: 700:0.0.2.188:700
    Target-attachment-id      Type  St      Time last up      # Up trans
    800:0.0.3.32:800          rmt   Up      Sep 18 01:17:38 2013      1
    Remote PE: 10.255.10.1, Negotiated control-word: Yes (Null), Encapsulation:
ETHERNET
    Incoming label: 300016, Outgoing label: 300048
    Negotiated PW status TLV: Yes
    local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
  Local source-attachment-id: 800:0.0.3.32:800
    Target-attachment-id      Type  St      Time last up      # Up trans
    700:0.0.2.188:700          rmt   Up      Sep 18 01:17:38 2013      1
    Remote PE: 10.255.3.1, Negotiated control-word: Yes (Null), Encapsulation:
ETHERNET
    Incoming label: 300000, Outgoing label: 300064
    Negotiated PW status TLV: Yes
    local PW status code: 0x00000000, Neighbor PW status code: 0x00000000
  Pseudowire Switching Points :
    Local address      Remote address      Status
    10.255.3.1         10.255.14.1        forwarding

```

Check the following fields in the output to verify that MS-PW is established between the T-PE devices:

- **Target-attachment-id**—Check if the TAI value is the SAI value of T-PE2.
- **Remote PE**—Check if the T-PE1 and S-PE2 loopback addresses are listed.
- **Negotiated PW status TLV**—Ensure that the value is **Yes**.
- **Pseudowire Switching Points**—Check if the switching points are listed from S-PE2 to T-PE2.

Meaning

MS-PW is established between T-PE1 and T-PE2 in the forwarding direction.

Checking the MS-PW Connections on S-PE2

Purpose

Make sure that all of the FEC 129 MS-PW connections come up correctly for the mspw routing instance.

Action

From operational mode, run the **show l2vpn connections instance __MSPW__ extensive** command.

```
user@S-PE2> show l2vpn connections instance __MSPW__ extensive
```

Layer-2 VPN connections:

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unknown connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy
RS -- remote site standby	SN -- Static Neighbor
LB -- Local site not best-site	RB -- Remote site not best-site
VM -- VLAN ID mismatch	

Legend for interface status

Up -- operational
Dn -- down

Instance: __MSPW__

L2vpn-id: 100:15

Local source-attachment-id: 700:0.0.2.188:700

Target-attachment-id	Type	St	Time last up	# Up trans
800:0.0.3.32:800	rmt	Up	Sep 18 00:58:55 2013	1

Remote PE: 10.255.2.1, Negotiated control-word: Yes (Null), Encapsulation:

ETHERNET

Incoming label: 300064, Outgoing label: 300000

Negotiated PW status TLV: Yes

local PW status code: 0x00000000, Neighbor PW status code: 0x00000000

Pseudowire Switching Points :

Local address	Remote address	Status
10.255.2.1	10.255.10.1	forwarding


```

Local source-attachment-id: 800:0.0.3.32:800
  Target-attachment-id      Type  St      Time last up      # Up trans
  700:0.0.2.188:700         rmt   Up      Sep 18 00:58:55 2013      1
    Remote PE: 10.255.14.1, Negotiated control-word: Yes (Null), Encapsulation:
ETHERNET
    Incoming label: 300048, Outgoing label: 300112
    Negotiated PW status TLV: Yes
    local PW status code: 0x00000000, Neighbor PW status code: 0x00000000

```

Check the following fields in the output to verify that MS-PW is established between the T-PE devices:

- **Target-attachment-id**—Check if the TAI value is the SAI value of T-PE1.
- **Remote PE**—Check if the S-PE1 and T-PE2 loopback addresses are listed.
- **Negotiated PW status TLV**—Ensure that the value is **Yes**.
- **Pseudowire Switching Points**—Check if the switching points are listed from S-PE1 to T-PE1.

Meaning

MS-PW is established between T-PE1 and T-PE2 in the reverse direction.

Checking the MS-PW Connections on T-PE2

Purpose

Make sure that all of the FEC 129 MS-PW connections come up correctly.

Action

From operational mode, run the **show l2vpn connections extensive** command.

```
user@T-PE2> show l2vpn connections extensive
```

```
Layer-2 VPN connections:
```

```
Legend for connection status (St)
```

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision


```

LN -- local site not designated  LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status  IL -- no incoming label
MM -- MTU mismatch                MI -- Mesh-Group ID not available
BK -- Backup connection           ST -- Standby connection
PF -- Profile parse failure        PB -- Profile busy
RS -- remote site standby          SN -- Static Neighbor
LB -- Local site not best-site     RB -- Remote site not best-site
VM -- VLAN ID mismatch

```

Legend for interface status

Up -- operational

Dn -- down

Instance: ms-pw

L2vpn-id: 100:15

Number of local interfaces: 1

Number of local interfaces up: 1

ge-2/0/0.0

Local source-attachment-id: 700:0.0.2.188:700 (CE2)

Target-attachment-id	Type	St	Time last up	# Up trans
800:0.0.3.32:800	rmt	Up	Sep 18 01:35:21 2013	1

Remote PE: 10.255.3.1, Negotiated control-word: Yes (Null)

Incoming label: 300112, Outgoing label: 300048

Negotiated PW status TLV: Yes

local PW status code: 0x00000000, Neighbor PW status code: 0x00000000

Local interface: ge-2/0/0.0, Status: Up, Encapsulation: ETHERNET

Pseudowire Switching Points :

Local address	Remote address	Status
10.255.3.1	10.255.2.1	forwarding
10.255.2.1	10.255.10.1	forwarding

Connection History:

Sep 18 01:35:21 2013	status update timer	
Sep 18 01:35:21 2013	PE route changed	
Sep 18 01:35:21 2013	Out lbl Update	300048
Sep 18 01:35:21 2013	In lbl Update	300112
Sep 18 01:35:21 2013	loc intf up	ge-2/0/0.0

Check the following fields in the output to verify that MS-PW is established between the T-PE devices:

- **Target-attachment-id**—Check if the TAI value is the SAI value of T-PE1.
- **Remote PE**—Check if the T-PE1 loopback address is listed.

- **Negotiated PW status TLV**—Ensure that the value is **Yes**.
- **Pseudowire Switching Points**—Check if the switching points are listed from S-PE2 to S-PE1 and from S-PE1 to T-PE1.

Meaning

MS-PW is established between T-PE1 and T-PE2 in the reverse direction.

Troubleshooting

IN THIS SECTION

- [Ping | 1297](#)
- [Bidirectional Forwarding Detection | 1298](#)
- [Traceroute | 1299](#)

To troubleshoot the MS-PW connection, see:

Ping

Problem

How to check the connectivity between the T-PE devices and between a T-PE device and an intermediary device.

Solution

Verify that T-PE1 can ping T-PE2. The **ping mpls l2vpn fec129** command accepts SAs and TAs as integers or IP addresses and also allows you to use the CE-facing interface instead of the other parameters (**instance**, **local-id**, **remote-id**, **remote-pe-address**).

Checking Connectivity Between T-PE1 and T-PE2

```
user@T-PE1> ping mpls l2vpn fec129 instance FEC129-VPWS local-id 800:800:800 remote-pe-address
10.255.14.1 remote-id 700:700:700
```

```
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

```
user@T-PE1> ping mpls l2vpn fec129 interface ge-3/1/2
```



```

!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

```

Checking Connectivity Between T-PE1 and S-PE2

user@T-PE1> **ping mpls l2vpn fec129 interface ge-3/1/2 bottom-label-ttl 2**

```

!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

```

Bidirectional Forwarding Detection

Problem

How to use BFD to troubleshoot the MS-PW connection from the T-PE device.

Solution

From operational mode, verify the **show bfd session extensive** command output.

user@T-PE1> **show bfd session extensive**

```

Address                State      Interface      Detect   Transmit
Time              Interval Multiplier
198.51.100.7          Up        ge-3/1/0.0     0.900    0.300      3
Client FEC129-OAM, TX interval 0.300, RX interval 0.300
Session up time 03:12:42
Local diagnostic None, remote diagnostic None
Remote state Up, version 1
Replicated
Session type: VCCV BFD
Min async interval 0.300, min slow interval 1.000
Adaptive async TX interval 0.300, RX interval 0.300
Local min TX interval 0.300, minimum RX interval 0.300, multiplier 3
Remote min TX interval 0.300, min RX interval 0.300, multiplier 3
Local discriminator 19, remote discriminator 19
Echo mode disabled/inactive
Remote is control-plane independent
L2vpn-id 100:15, Local-id 800:0.0.3.32:800, Remote-id 700:0.0.2.188:700
Session ID: 0x103

```



```
1 sessions, 1 clients
Cumulative transmit rate 3.3 pps, cumulative receive rate 3.3 pps
```

Traceroute

Problem

How to verify that MS-PW was established.

Solution

From operational mode, verify **traceroute** output.

```
user@T-PE1> traceroute mpls l2vpn fec129 interface interface
```

```
Probe options: ttl 64, retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1		FEC129	10.255.10.1	(null)	Success
2		FEC129	10.255.2.1	10.255.10.1	Success
3		FEC129	10.255.3.1	10.255.2.1	Success
4		FEC129	10.255.14.1	10.255.2.1	Egress

```
Path 1 via ge-3/1/2 destination 198.51.100.0
```

RELATED DOCUMENTATION

[Understanding Multisegment Pseudowire for FEC 129 | 1248](#)

MPLS Stitching For Virtual Machine Connection

IN THIS SECTION

- [When Would I Use Stitching? | 1300](#)
- [How Does MPLS Stitching Work? | 1300](#)
- [How Do I Configure Stitching? | 1301](#)


- Which Switches Support Stitching? | 1301
- Q&A | 1302

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 either BGP labeled unicast, RSVP or LDP, to provide the routes dictated by the labels.

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 104 on page 1300](#).

Figure 104: Virtual Machines on Either Side of Routers



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 104 on page 1300](#). 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.

Traffic arriving at ToR1 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 I-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 forwards 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** has been added under **transit** to resolve the remote next-hop. For example, instead of **set protocols mpls static-label-switched-path static-to-ToR2 transit 1000000 next-hop 10.9.82.47**, a top-of-rack switch redirects packets to another top-of-rack switch with **set protocols mpls static-label-switched-path static-to-ToR2 transit 1000000 stitch**. 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.

See the complete example for stitching at *Using MPLS Stitching with BGP to Connect Virtual Machines* for more information.

Which Switches Support Stitching?

See [Feature Explorer](#) for the list of switches that support the [MPLS Stitching For Virtual Machine Connections](#) 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.

RELATED DOCUMENTATION

[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)

Using MPLS Stitching with BGP to Connect Virtual Machines

TDM Pseudowires Overview

A TDM pseudowire acts as Layer 2 circuit or service for T1 and E1 circuit signals across an MPLS packet-switched network. On ACX Series routers, you configure a TDM pseudowire with Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP) on the ACX Series built-in channelized T1 and E1 interfaces. When you configure a TDM pseudowire, the network between the customer edge (CE) routers appears transparent to the CE routers, making it seem that the CE routers are directly connected. With the SAToP configuration on the provider edge (PE) router's T1 and E1 interfaces, the interworking function (IWF) forms a payload (frame) that contains the CE router's T1 and E1 Layer 1 data and control word. This data is transported to the remote PE over the pseudowire. The remote PE removes all the Layer 2 and MPLS headers added in the network cloud and forwards the control word and the Layer 1 data to the remote IWF, which in turn forwards the data to the remote CE router.

RELATED DOCUMENTATION

Understanding Encapsulation on an Interface

SAToP Emulation on T1 and E1 Interfaces Overview

Configuring SAToP Emulation on Channelized T1 and E1 Interfaces

[Pseudowire Overview for ACX Series Universal Metro Routers | 1246](#)

ATM Pseudowire Overview

[Ethernet Pseudowire Overview | 1241](#)

Example: TDM Pseudowire Base Configuration

IN THIS SECTION

- [Requirements | 1303](#)
- [Overview of a TDM Pseudowire Base Configuration | 1303](#)
- [Configuring an TDM Pseudowire | 1303](#)

Requirements

The following is a list of the hardware and software requirements for this configuration.

- One ACX Series router
- Junos OS Release 12.2 or later

Overview of a TDM Pseudowire Base Configuration

The configuration shown here is the base configuration of an TDM pseudowire with T1 framing on an ACX Series router. This configuration is for one provider edge router. To complete the TDM pseudowire configuration, you need to repeat this configuration on an other provider edge router in the Multiprotocol Label Switched (MPLS) network.

Configuring an TDM Pseudowire

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in 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:

```
set chassis fpc 0 pic 0 framing t1
set interfaces ct1-0/0/0 no-partition interface-type t1
set interfaces t1-0/0/0 encapsulation satop
set interfaces t1-0/0/0 unit 0
set interfaces ge-0/2/0 unit 0 family inet address 20.1.1.2/24
set interfaces ge-0/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 70.1.1.1/32
set protocols rsvp interface ge-0/2/0.0
```



```

set protocols mpls no-cspf
set protocols mpls label-switched-path PE1-to-PE2 to 40.1.1.1
set protocols mpls interface ge-0/2/0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-0/2/0.0
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 40.1.1.1 interface t1-0/0/0.0 virtual-circuit-id
1

```

NOTE: To configure a TDM pseudowire with E1 framing, include the **e1** statement at the **[edit chassis fpc 0 pic 0 framing]** hierarchy level instead of the **t1** statement shown in this example.

Step-by-Step Procedure

1. Configure the framing format:

```

[edit]
user@host# edit chassis
[edit chassis]
user@host# set fpc 0 pic 0 framing t1

```

2. Create a T1 interface on a channelized T1 interface (**ct1**) and enable full channelization with the **no-partition** statement. On the logical T1 interface, set the Structure-Agnostic TDM over Packet (SAToP) encapsulation mode.

```

[edit]
user@host# edit interfaces
[edit interfaces]
user@host# set ct1-0/0/0 no-partition interface-type t1
user@host# set t1-0/0/0 encapsulation satop
user@host# set t1-0/0/0 unit 0

```


3. Create a Gigabit Ethernet interface and enable MPLS on that interface. Create the loopback (**lo0**) interface:

```
[edit interfaces]
user@host# set ge-0/2/0 unit 0 family inet address 20.1.1.2/24
user@host# set ge-0/2/0 unit 0 family mpls
user@host# set lo0 unit 0 family inet address 70.1.1.1/32
```

4. Enable the MPLS and RSVP protocols on the MPLS interface—**ge-0/2/0.0**:

```
[edit]
user@host# edit protocols
[edit protocols]
user@host# set rsvp interface ge-0/2/0.0
user@host# set mpls interface ge-0/2/0.0
```

5. Configure LDP. If you configure RSVP for a pseudowire, you must also configure LDP:

```
[edit protocols]
user@host# set ldp interface ge-0/2/0.0
user@host# set ldp interface lo0.0
```

6. Configure a point-to-point label-switched path (LSP) and disable constrained-path LSP computation:

```
[edit protocols]
user@host# set mpls label-switched-path PE1-to-PE2 to 40.1.1.1
user@host# set mpls no-cspf
```

7. Configure OSPF and enable traffic engineering on the MPLS interface—**ge-0/2/0.0**, and on the loopback (**lo0**) interface:

```
[edit protocols]
user@host# set ospf traffic-engineering
user@host# set ospf area 0.0.0.0 interface ge-0/2/0.0
user@host# set ospf area 0.0.0.0 interface lo0.0 passive
```


8. Uniquely identify a Layer 2 circuit for the TDM pseudowire:

```
[edit protocols]
user@host# set l2circuit neighbor 40.1.1.1 interface t1-0/0/0.0 virtual-circuit-id 1
```

Results

```
[edit]
user@host# show
chassis {
    fpc 0 {
        pic 0 {
            framing t1;
        }
    }
}
interfaces {
    ct1-0/0/0 {
        no-partition interface-type t1;
    }
    t1-0/0/0 {
        encapsulation satop;
        unit 0;
    }
    ge-0/2/0 {
        unit 0 {
            family inet {
                address 20.1.1.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 70.1.1.1/32;
            }
        }
    }
}
protocols {
    rsvp {
        interface ge-0/2/0.0;
```



```

}
mpls {
    no-cspf;
    label-switched-path PE1-to-PE2 {
        to 40.1.1.1;
    }
    interface ge-0/2/0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-0/2/0.0;
        interface lo0.0 {
            passive;
        }
    }
}
ldp {
    interface ge-0/2/0.0;
    interface lo0.0;
}
l2circuit {
    neighbor 40.1.1.1 {
        interface t1-0/0/0.0 {
            virtual-circuit-id 1;
        }
    }
}
}

```

RELATED DOCUMENTATION

[Pseudowire Overview for ACX Series Universal Metro Routers](#) | 1246

[TDM Pseudowires Overview](#) | 1302

Configuring Load Balancing for Ethernet Pseudowires

You can configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires. You can also configure load balancing for Ethernet pseudowires based on IP information. The option to include IP information in the hash key provides support for Ethernet circuit cross-connect (CCC) connections.

NOTE: This feature is supported only on M120, M320, MX Series, and T Series routers.

To configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires, include the **ether-pseudowire** statement at the **[edit forwarding-options hash-key family mpls payload]** hierarchy level:

```
[edit forwarding-options]
hash-key {
  family mpls {
    (label-1 | no-labels);
    payload {
      ether-pseudowire;
    }
  }
}
```

NOTE: You must also configure either the **label-1** or the **no-labels** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level.

You can also configure load balancing for Ethernet pseudowires based on IP information. This functionality provides support for load balancing for Ethernet cross-circuit connect (CCC) connections. To include IP information in the hash key, include the **ip** statement at the **[edit forwarding-options hash-key family mpls payload]** hierarchy level:

```
[edit forwarding-options]
hash-key {
  family mpls {
    (label-1 | no-labels);
    payload {
      ip;
    }
  }
}
```


NOTE: You must also configure either the **label-1** or **no-labels** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level.

You can configure load balancing for IPv4 traffic over Ethernet pseudowires to include only Layer 3 IP information in the hash key. To include only Layer 3 IP information, include the **layer-3-only** option at the **[edit forwarding-options family mpls hash-key payload ip]** hierarchy level:

```
[edit forwarding-options]
hash-key {
  family mpls {
    (label-1 | no-labels);
    payload {
      ip {
        layer-3-only;
      }
    }
  }
}
```

NOTE: You must also configure either the **label-1** or **no-labels** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level.

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include the **family multiservice** statement at the **[edit forwarding-options hash-key]** hierarchy level:

```
family multiservice {
  destination-mac;
  source-mac;
}
```

To include the destination-address MAC information in the hash key, include the **destination-mac** option. To include the source-address MAC information in the hash key, include the **source-mac** option.

NOTE: Any packets that have the same source and destination address will be sent over the same path.

NOTE: You can configure per-packet load balancing to optimize VPLS traffic flows across multiple paths.

NOTE: Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.

NOTE: ACX Series routers do not support VPLS.

RELATED DOCUMENTATION

| *Junos OS VPNs Library for Routing Devices*

Configuring Class-of-Service (CoS) for MPLS

IN THIS CHAPTER

- [Configuring Class of Service for MPLS LSPs | 1311](#)
- [Configuring MPLS Rewrite Rules | 1315](#)
- [Configuring CoS Bits for an MPLS Network \(CLI Procedure\) | 1317](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)
- [Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) | 1323](#)
- [Understanding Using CoS with MPLS Networks on EX Series Switches | 1324](#)
- [Example: Combining CoS with MPLS on EX Series Switches | 1329](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules | 1345](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers | 1349](#)
- [Configuring CoS Bits for an MPLS Network | 1351](#)
- [Configuring a Global MPLS EXP Classifier | 1352](#)

Configuring Class of Service for MPLS LSPs

IN THIS SECTION

- [Class of Service for MPLS Overview | 1312](#)
- [Configuring the MPLS CoS Values | 1312](#)
- [Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value | 1314](#)

The following sections provide an overview of MPLS class of service (CoS) and describe how to configure the MPLS CoS value:

Class of Service for MPLS Overview

When IP traffic enters an LSP tunnel, the ingress router marks all packets with a CoS value, which is used to place the traffic into a transmission priority queue. On the router, for SDH/SONET and T3 interfaces, each interface has four transmit queues. The CoS value is encoded as part of the MPLS header and remains in the packets until the MPLS header is removed when the packets exit from the egress router. The routers within the LSP utilize the CoS value set at the ingress router. The CoS value is encoded by means of the CoS bits (also known as the EXP or experimental bits). For more information, see [“MPLS Label Allocation” on page 420](#).

MPLS class of service works in conjunction with the router’s general CoS functionality. If you do not configure any CoS features, the default general CoS settings are used. For MPLS class of service, you might want to prioritize how the transmit queues are serviced by configuring weighted round-robin, and to configure congestion avoidance using random early detection (RED)..

Configuring the MPLS CoS Values

When traffic enters an LSP tunnel, the CoS value in the MPLS header is set in one of three 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. *Default MPLS EXP Classifier* explains the default MPLS CoS values, and summarizes how the CoS values 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.
- You set an MPLS EXP rewrite rule to override the default behavior.

To set a fixed CoS value on all packets entering the LSP, include the **class-of-service** statement:

```
class-of-service cos-value;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit protocols mpls label-switched-path *path-name*]
- [edit protocols mpls label-switched-path *path-name* primary *path-name*]
- [edit protocols mpls label-switched-path *path-name* secondary *path-name*]
- [edit protocols rsvp interface *interface-name* link-protection]
- [edit protocols rsvp interface *interface-name* link-protection bypass *destination*]
- [edit logical-systems *logical-system-name* protocols mpls]

- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name* primary *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name* secondary *path-name*]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection bypass *destination*]

The CoS value set using the **class-of-service** statement at the [edit protocols mpls] hierarchy level supersedes the CoS value set at the [edit class-of-service] hierarchy level for an interface. Effectively, the CoS value configured for an LSP overrides the CoS value set for an interface.

The **class-of-service** statement at the [edit protocols mpls label-switched-path] hierarchy level assigns an initial EXP value for the MPLS shim header of packets in the LSP. This value is initialized at the ingress routing device only and overrides the rewrite configuration established for that forwarding class. However, the CoS processing (weighted round robin [WRR] and RED) of packets entering the ingress routing device is not changed by the **class-of-service** statement on an MPLS LSP. Classification is still based on the behavior aggregate (BA) classifier at the [edit class-of-service] hierarchy level or the multifield classifier at the [edit firewall] hierarchy level.

BEST PRACTICE: We recommend configuring all routing devices along the LSP to have the same input classifier for EXP, and, if a rewrite rule is configured, all routing devices should have the same rewrite configuration. Otherwise, traffic at the next LSR might be classified into a different forwarding class, resulting in a different EXP value being written to the EXP header.

The CoS value can be a decimal number from 0 through 7. This number corresponds to a 3-bit binary number. The high-order 2 bits of the CoS value select which transmit queue to use on the outbound interface card.

The low-order bit of the CoS value is treated as the PLP bit and is used to select the RED drop profile to use on the output queue. If the low-order bit is 0, the non-PLP drop profile is used, and if the low-order bit is 1, the PLP drop profile is used. It is generally expected that RED will more aggressively drop packets that have the PLP bit set. For more information about RED and drop profiles, see *Managing Congestion Using RED Drop Profiles and Packet Loss Priorities*.

NOTE: Configuring the PLP drop profile to drop packets more aggressively (for example, setting the CoS value from 6 to 7) decreases the likelihood of traffic getting through.

Table 28 on page 1314 summarizes how MPLS CoS values correspond to the transmit queue and PLP bit. Note that in MPLS, the mapping between the CoS bit value and the output queue is hard-coded. You cannot configure the mapping for MPLS; you can configure it only for IPv4 traffic flows, as described in *Understanding How Forwarding Classes Assign Classes to Output Queues*.

Table 28: MPLS CoS Values

MPLS CoS Value	Bits	Transmit Queue	PLP Bit
0	000	0	Not set
1	001	0	Set
2	010	1	Not set
3	011	1	Set
4	100	2	Not set
5	101	2	Set
6	110	3	Not set
7	111	3	Set

Because the CoS value is part of the MPLS header, the value is associated with the packets only as they travel through the LSP tunnel. The value is not copied back to the IP header when the packets exit from the LSP tunnel.

To configure class of service (CoS) for Multiprotocol Label Switching (MPLS) packets in a label-switched path (LSP):

1. Specify the CoS value

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.

Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value

For Ethernet interfaces installed on a T Series router or an M320 router with a peer connection to an M Series router or a T Series router, you can rewrite both MPLS CoS and IEEE 802.1p values to a configured value (the MPLS CoS values are also known as the EXP or experimental bits). Rewriting these values allows you to pass the configured value to the Layer 2 VLAN path. To rewrite both the MPLS CoS and IEEE 802.1p values, you must include the EXP and IEEE 802.1p rewrite rules in the class of service interface configuration. The EXP rewrite table is applied when you configure the IEEE 802.1p and EXP rewrite rules.

For information about how to configure the EXP and IEEE 802.1p rewrite rules, see *Rewriting Packet Headers to Ensure Forwarding Behavior*.

RELATED DOCUMENTATION

| *Default MPLS EXP Classifier*

Configuring MPLS Rewrite Rules

IN THIS SECTION

- [Rewriting the EXP Bits of All Three Labels of an Outgoing Packet | 1315](#)
- [Rewriting MPLS and IPv4 Packet Headers | 1316](#)

You can apply a number of different rewrite rules to MPLS packets.

For more information about how to configure statements at the **[edit class-of-service]** hierarchy level, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

The following sections describe how you can apply rewrite rules to MPLS packets:

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

In interprovider, carrier-of-carrier, and complex traffic engineering scenarios, it is sometimes necessary to push three labels on the next hop.

By default, on M Series routers except the M320, the top MPLS EXP label of an outgoing packet is not rewritten when you configure swap-push-push and triple-push operations. You can rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining the class of service (CoS) of an incoming MPLS or non-MPLS packet.

To push three labels on incoming MPLS packets, include the **exp-swap-push-push default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
exp-swap-push-push default;
```


To push three labels on incoming non-MPLS packets, include the **exp-push-push-push default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
exp-push-push-push default;
```

For more information about how to configure statements at the **[edit class-of-service]** hierarchy level, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

Rewriting MPLS and IPv4 Packet Headers

You can apply a rewrite rule to MPLS and IPv4 packet headers simultaneously. This allows you to initialize MPLS EXP and IP precedence bits at LSP ingress. You can configure different rewrite rules depending on whether the traffic is VPN or non-VPN.

To rewrite MPLS and IPv4 packet headers, include the **protocol** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules exp *rewrite-rule-name*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name]
protocol types;
```

Use the **protocol** statement to specify the types of MPLS packets and packet headers to which to apply the rewrite rule. The MPLS packet can be a standard MPLS packet or an MPLS packet with an IPv4 payload. Specify the type of MPLS packet by using the following options:

- **mpls-any**—Applies the rewrite rule to MPLS packets and writes the code point value to MPLS headers.
- **mpls-inet-both**—Applies the rewrite rule to VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T Series (except T4000 routers) and M320 routers. On M Series routers, except the M320, the **mpls-inet-both** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.
- **mpls-inet-both-non-vpn**—Applies the rewrite rule to any non-VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T Series and M320 routers. On M Series routers, except the M320, the **mpls-inet-both-non-vpn** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.

For a detailed example on how to configure rewrite rules for MPLS and IPv4 packets and for more information about how to configure class of service, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

Configuring CoS Bits for an MPLS Network (CLI Procedure)

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 [Junos OS Class of Service Configuration Guide](#) explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

To set a fixed CoS value on all packets entering the LSP:

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

Example: Configuring CoS on EX Series Switches

[Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)

[Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)

[Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) | 1323](#)

Defining CoS Rewrite Rules (CLI Procedure)

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on a provider edge (PE) switch that is using IP Over MPLS.

This task describes how to create a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE switch. It includes configuring a policer firewall filter and applying it to the customer-edge interface of the ingress PE switch. The policer firewall filter ensures that the amount of traffic forwarded through the MPLS tunnel never exceeds the requested bandwidth allocation.

Before you begin, configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 95.](#)
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99.](#)

This topic includes:

1. [Configuring CoS | 1318](#)
2. [Configuring an LSP Policer | 1319](#)

Configuring CoS

To configure CoS on a provider edge switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers dscp classifier-name import default
```

2. Add a forwarding class to this custom DSCP classifier and specify a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers dscp classifier-name forwarding-class forwarding-class loss-priority
loss-priority code-points code-point
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
```



```
user@switch# set rewrite-rules exp e1 forwarding-class forwarding-class loss-priority loss-priority
code-points code-point
```

4. On EX8200 switches only, bind the custom EXP rewrite rule to the interface:

```
[edit class-of-service]
user@switch# set class-of-service interfaces interface unit unit rewrite-rules exp e1
```

Configuring an LSP Policer

To configure an LSP policer:

NOTE: You cannot configure LSP policers on EX8200 switches. EX8200 switches do not support LSP policers.

1. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the customer-edge-interface:

```
[edit firewall]
user@switch# set policer mypolicer if-exceeding bandwidth-limit 500m
```

2. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

```
[edit firewall policer]
user@switch# set mypolicer if-exceeding burst-size-limit 33553920
```

3. Discard traffic that exceeds the rate limits for this policer:

```
[edit firewall policer]
user@switch# set mypolicer then discard
```

4. To reference the policer, configure a filter term that includes the policer action:

```
[edit firewall]
user@switch# set family inet filter myfilter term t1 then policer mypolicer
```

5. Apply the filter to the customer-edge interface:

```
[edit interfaces]
user@switch# set ge-2/0/3 unit 0 family inet address 192.168.121.1/16 policing filter myfilter
```


NOTE: You can also configure schedulers and shapers as needed. See *Defining CoS Schedulers and Scheduler Maps (CLI Procedure)*.

RELATED DOCUMENTATION

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

Assigning CoS Components to Interfaces (CLI Procedure)

Configuring Policers to Control Traffic Rates (CLI Procedure)

Understanding the Use of Policers in Firewall Filters

Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on a provider edge (PE) switch that is using MPLS over circuit-cross connect (CCC).

NOTE: On EX Series switches other than EX8200 switches, if you are using MPLS over CCC, you can use only one DSCP or IP precedence classifier and only one IEEE 802.1p classifier on the CCC interfaces.

This procedure is for creating a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE. It also includes enabling a policer on the label-switched path (LSP) of the ingress PE to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

This topic includes:

1. [Configuring CoS | 1321](#)
2. [Configuring an LSP Policer | 1322](#)

Configuring CoS

To configure CoS on a provider edge switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers dscp classifier-name import default
```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers dscp classifier-name forwarding-class forwarding-class loss-priority loss-priority code-points code-point
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class forwarding-class loss-priority loss-priority code-point code-point
```

4. Bind the DSCP classifier to the CCC interface:

```
[edit ]
user@switch# set class-of-service interfaces interface unit unit classifier classifier-name
```

5. On EX8200 switches only, bind the custom EXP rewrite rule to the interface:

```
[edit class-of-service]
user@switch# set class-of-service interfaces interface unit unit rewrite-rules exp e1
```


Configuring an LSP Policer

To configure an LSP policer:

NOTE: You cannot configure LSP policers on EX8200 switches. EX8200 switches do not support LSP policers.

1. Specify the number of bits per second permitted, on average, for the policer, which will later be applied to the LSP:

```
[edit firewall]  
set policer mypolicer if-exceeding bandwidth-limit 500m
```

2. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

```
[edit firewall policer]  
set mypolicer if-exceeding burst-size-limit 33553920
```

3. Discard traffic that exceeds the rate limits for this policer:

```
[edit firewall policer]  
set mypolicer then discard
```

4. To reference the policer, configure a filter term that includes the policer action:

```
[edit firewall]  
user@switch# set family any filter myfilter term t1 then policer mypolicer
```

5. Apply the filter to the LSP:

```
[edit protocols mpls]  
set label-switched-path lsp_to_pe2_ge1 policing filter myfilter
```

NOTE: You can also configure schedulers and shapers as needed. See *Defining CoS Schedulers and Scheduler Maps (CLI Procedure)*.

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Assigning CoS Components to Interfaces \(CLI Procedure\)](#)

[Configuring Policers to Control Traffic Rates \(CLI Procedure\)](#)

[Understanding the Use of Policers in Firewall Filters](#)

Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure)

You can add class-of-service (CoS) components to your MPLS networks on EX Series switches to achieve end-to-end Differentiated Services to match your specific business requirements. The configuration of CoS components on the provider switches is the same regardless of whether the provider edge (PE) switches are using MPLS over CCC or IP over MPLS.

This task shows how to configure a custom EXP classifier and custom EXP rewrite rule on the provider switch.

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

4. On EX8200 switches only, bind the custom EXP rewrite rule to the interface:

```
[edit class-of-service]
user@switch# set class-of-service interfaces ge-0/0/2 unit 0 rewrite-rules exp e1
```


NOTE: You can also configure schedulers and shapers as needed. See *Defining CoS Schedulers and Scheduler Maps (CLI Procedure)*.

RELATED DOCUMENTATION

| [Example: Configuring CoS on EX Series Switches](#)

Understanding Using CoS with MPLS Networks on EX Series Switches

IN THIS SECTION

- [EXP Classifiers and EXP rewrite Rules | 1325](#)
- [Guidelines for Using CoS Classifiers on CCCs | 1325](#)
- [Using CoS Classifiers with IP over MPLS | 1326](#)
- [Setting CoS Bits in an MPLS Header | 1326](#)
- [EXP Rewrite Rules | 1328](#)
- [Policer | 1328](#)
- [Schedulers | 1329](#)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. See *EX Series Switch Software Features Overview* for a complete list of the Junos OS MPLS features that are supported on specific EX Series switches.

Juniper Networks EX Series Ethernet Switches support Differentiated Service Code Point (DSCP) or IP precedence and IEEE 802.1p CoS classifiers on the customer-edge interfaces of the ingress provider edge (PE) switch. DSCP or IP precedence classifiers are used for Layer 3 packets. IEEE 802.1p is used for Layer 2 packets.

When a packet enters a customer-edge interface of the ingress PE switch, the switch associates the packet with a particular CoS servicing level before putting the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch. The CoS value that was embedded in the classifier is translated and encoded in the MPLS header by means of the EXP or experimental bits. EX Series switches enable a default EXP classifier and a default EXP rewrite rule. For more information about EXP classifiers and EXP rewrite rules, see *EXP Classifiers and EXP rewrite Rules*.

This topic includes:

EXP Classifiers and EXP rewrite Rules

EX Series switches enable a default EXP classifier and a default EXP rewrite rule. You can configure a custom EXP classifier and a custom EXP rewrite rule if you prefer. However, the switch supports only one type of EXP classifier (default or custom) and only one EXP rewrite rule (default or custom).

You do not bind the EXP classifier or the EXP rewrite rule to individual interfaces. The switch automatically and implicitly applies the default or the custom EXP classifier and the default or the custom EXP rewrite rule to the appropriate MPLS-enabled interfaces. Because rewrite rules affect only egress interfaces, the switch applies the EXP rewrite rule only to those MPLS interfaces that are transmitting MPLS packets (not to the MPLS interfaces that are receiving the packets).

After traversing the MPLS tunnel, the traffic flows out from the egress provider edge (PE) switch. Before the traffic leaves the egress interface, the egress PE switch copies the EXP bits from the MPLS header to the most significant bits in the original IP packet--- that is, to the IP precedence bits. Note that this is the default behavior only on Juniper Networks EX8200 Ethernet Switches (standalone or Virtual Chassis) that are configured for MPLS.

Guidelines for Using CoS Classifiers on CCCs

When you are configuring CoS for MPLS over circuit cross-connect (CCC), there are some additional guidelines, as follows:

- You *must* explicitly bind a CoS classifier to the CCC interface on the ingress PE switch.
- You *must* use the same DSCP, IP precedence, or IEEE 802.1p classifier on CCC interfaces. However, if the CCC interfaces are on the same switch, you cannot configure both a DSCP and an IP precedence

classifier on these interfaces. Thus, if you configure one CCC interface to use a DSCP classifier DSCP1, you cannot configure another CCC interface to use another DSCP classifier DSCP2. All the CCC interfaces on the switch must use the same DSCP (or IP precedence) classifier and the same IEEE 802.1p classifier.

- You *cannot* configure one CCC interface to use a DSCP classifier and another CCC interface to use an IP precedence classifier, because these classifier types overlap.
- You *can* configure one CCC interface to use a DSCP classifier and another CCC interface to use IEEE 802.1p classifier.
- You *can* configure one CCC interface to use both a DSCP and an IEEE 802.1p classifier. If you configure a CCC interface to use both these classifiers, the DSCP classifier is used for routing Layer 3 packets and the IEEE 802.1p classifier is used for routing Layer 2 packets.
- You *can* configure one CCC interface to use both an IP precedence and an IEEE 802.1p classifier. If you configure a CCC interface to use both these classifiers, the IP precedence classifier is used for routing Layer 3 packets and the IEEE 802.1p classifier is used for routing Layer 2 packets.

NOTE: These guidelines are not applicable to Juniper Networks EX8200 Ethernet Switches (standalone or Virtual Chassis).

You can define multiple DSCP, IP precedence, and IEEE 802.1p classifiers for the non-CCC interfaces on a switch.

Using CoS Classifiers with IP over MPLS

When you are configuring CoS for IP over MPLS, the customer-edge interface uses the CoS configuration for the switch as the default. You do not have to bind a classifier to the customer-edge interface in this case. There are no restrictions on using multiple DSCP, IP precedence, and IEEE 802.1p classifiers on the same switch.

- You can modify the CoS classifier for a particular interface, but it is not required.
- You can configure a DSCP classifier, DSCP1 on the first interface, another DSCP classifier, DSCP2 on the second interface, and an IP precedence classifier on a third interface, and so forth.

Setting CoS Bits in an MPLS Header

When traffic enters an LSP tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet's CoS value. This behavior is the default, and no configuration is required. The *Class of Service User Guide (Routers and EX9200 Switches)* 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.

The CoS value can be a decimal number from 0 through 7. This number corresponds to a 3-bit binary number. The high-order 2 bits of the CoS value select which transmit queue to use on the outbound interface card.

The low-order bit of the CoS value is treated as the PLP bit and is used to select the RED drop profile to use on the output queue. If the low-order bit is 0, the non-PLP drop profile is used, and if the low-order bit is 1, the PLP drop profile is used. It is generally expected that random early detection (RED) will more aggressively drop packets that have the PLP bit set. For more information about RED and drop profiles, see the *Class of Service User Guide (Routers and EX9200 Switches)*.

NOTE: Configuring the PLP drop profile to drop packets more aggressively (for example, setting the CoS value from 6 to 7) decreases the likelihood of traffic getting through.

Table 29 on page 1327 summarizes how MPLS CoS values correspond to the transmit queue and PLP bit. Note that in MPLS, the mapping between the CoS bit value and the output queue is hard-coded. You cannot configure the mapping for MPLS; you can configure it only for IPv4 traffic flows, as described in the *Class of Service User Guide (Routers and EX9200 Switches)*.

Table 29: MPLS CoS Values

MPLS CoS Value	Bits	Transmit Queue	PLP Bit
0	000	0	Not set
1	001	0	Set
2	010	1	Not set
3	011	1	Set
4	100	2	Not set
5	101	2	Set
6	110	3	Not set
7	111	3	Set

Because the CoS value is part of the MPLS header, the value is associated with the packets only while they travel through the LSP tunnel. The value is not copied back to the IP header when the packets exit from the LSP tunnel.

NOTE: On EX8200 switches that run MPLS-based Layer 2 virtual private networks (VPNs):

- If you configure an LSP CoS, the EXP bits of the MPLS packet continue to use the same CoS values that are configured at the interface level.
- For Virtual Chassis, if the input and output interfaces are on different line cards, then the loss priority value that you configured on the first line card is not carried to the subsequent line cards. The loss priority for the outgoing traffic from the subsequent line cards is always set to low.

EXP Rewrite Rules

When traffic passes from the customer-edge interface to an MPLS interface, the DSCP, IP precedence, or IEEE 802.1p CoS classifier is translated into the EXP bits within the MPLS header. You cannot disable the default EXP rewrite rule, but you can configure your own custom EXP classifier and a custom EXP rewrite rule. You cannot bind the EXP classifier to individual MPLS interfaces; the switch applies it globally to all the MPLS-enabled interfaces on the switch.

Only one EXP rewrite rule (either default or custom) is supported on a switch. The switch applies it to all the egress interfaces on which MPLS is enabled.. This is, however, not the case with EX8200 switches. With EX8200 switches, you must explicitly apply the rewrite rule on each of the egress interfaces.

Policer

Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. During periods of congestion (when the total rate of queuing packets exceeds the rate of transmission), any new packets being sent to an interface can be dropped because there is no place to store them. You can configure a policer on the ingress PE switch to prevent this:

- If you are using MPLS over CCC, you bind the policer to the LSP. You cannot bind a policer to a CCC interface.
- If you are using IP over MPLS, you bind the policer to the **inet-family** customer-edge interface. You cannot bind a policer to the LSP when you are using IP over MPLS.

NOTE: You cannot configure LSP policers on EX8200 switches.

Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on EX Series switches. Default schedulers are provided for best-effort and network-control forwarding classes. If you are using assured-forwarding, expedited-forwarding, or any custom forwarding class, we recommend that you configure a scheduler to support that forwarding class. See *Understanding CoS Schedulers*.

RELATED DOCUMENTATION

Understanding CoS Classifiers

Example: Configuring CoS on EX Series Switches

[Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)

[Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)

[Configuring Rewrite Rules for EXP Classifiers on MPLS Networks \(CLI Procedure\)](#)

[Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) | 1323](#)

[Configuring CoS Bits for an MPLS Network \(CLI Procedure\) | 1317](#)

Example: Combining CoS with MPLS on EX Series Switches

IN THIS SECTION

- [Requirements | 1330](#)
- [Overview and Topology | 1330](#)
- [Configuring the Local PE Switch | 1333](#)
- [Configuring the Remote PE Switch | 1336](#)
- [Configuring the Provider Switch | 1337](#)
- [Verification | 1338](#)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. The CoS value is included within the MPLS label, which is passed through the network, enabling end-to-end CoS across the network.

MPLS services are often used to ensure better performance for low-latency applications such as VoIP and other business-critical functions. These applications place specific demands on a network for successful

transmission. CoS gives you the ability to control the mix of bandwidth, delay, jitter, and packet loss while taking advantage of the MPLS labeling mechanism.

This example shows how to configure CoS on an MPLS network that is using a unidirectional circuit cross-connect (CCC) from the ingress provider edge (PE) switch to the egress PE switch. for the customer-edge interface of the ingress provider edge (PE) switch. It describes adding the configuration of CoS components to the ingress PE switch, the egress PE switch, and the core provider switches of the existing MPLS network. Because of the unidirectional configuration, the DSCP classifier needs to be configured only on the ingress PE switch.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 10.1 or later for EX Series switches
- Three EX Series switches

Before you configure CoS with MPLS, be sure you have:

Configured an MPLS network with two PE switches and one provider switch. See [“Example: Configuring MPLS on EX8200 and EX4500 Switches” on page 59](#). This example assumes that an MPLS network has been configured using a cross circuit-connect (CCC).

Overview and Topology

This example describes adding custom classifiers and custom rewrite rules to switches in an MPLS network that is using MPLS over CCC.

It is a unidirectional configuration. Therefore, you need to configure custom classifiers and custom rewrite rules as follows:

- On the ingress PE switch: custom DSCP classifier and custom EXP rewrite rule
- On the egress PE switch: custom EXP classifier
- On the provider switch: customer EXP classifier and custom EXP rewrite rule

NOTE: You can also configure schedulers and shapers as needed. If you are using **assured-forwarding**, **expedited-forwarding**, or other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See *Defining CoS Schedulers and Scheduler Maps (CLI Procedure)*.

The example creates a custom DSCP classifier (**dscp1**) on the ingress PE switch and binds this classifier to the CCC interface. It includes configuration of a policer on the ingress PE switch. The policer is applied as

a filter on the label-switched path (LSP) **lsp_to_pe2_ge1**(created in [“Example: Configuring MPLS on EX8200 and EX4500 Switches” on page 59](#)) to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

This example creates a custom EXP rewrite rule (**exp1**) on the ingress PE switch, specifying a loss-priority and code point to be used for the expedited-forwarding class as the packet travels through the LSP. The switch applies this custom rewrite rule on the core interfaces **ge-0/0/5.0** and **ge-0/0/6.0**, which are the egress interfaces for this switch.

[Table 30 on page 1331](#) shows the CoS configuration components added to the ingress PE switch.

Table 30: CoS Configuration Components on the Ingress PE Switch

Property	Settings	Description
Local PE switch hardware	EX Series switch	PE-1
Policing filter configured and applied to the LSP.	policing filter mypolicer filter myfilter	Name of the rate-limiting policer. Name of the filter, which refers to the policer
Custom DSCP classifier	dscp1	Specifies the name of the custom DSCP classifier
Custom EXP rewrite rule	e1	Name of the custom EXP rewrite rule.
Customer-edge interface	ge-0/0/1.0	Interface that receives packets from devices outside the network. The custom DSCP classifier must be specified on this CCC interface.
Core interfaces	ge-0/0/5.0 and ge-0/0/6.0	Interfaces that transmit MPLS packets to other switches within the MPLS network. The EXP rewrite rule is applied implicitly to these interfaces.

[Table 31 on page 1332](#) shows the CoS configuration components added to the egress PE switch in this example.

Table 31: CoS Configuration Components of the Egress PE Switch

Property	Settings	Description
Remote provider edge switch hardware	EX Series switch	PE-2
Custom EXP classifier	exp1	Name of custom EXP classifier
Customer-edge interface	ge-0/0/1.0	Interface that transmits packets from this network to devices outside the network. No CoS classifier is specified for this interface. A scheduler can be specified.
Core interfaces	ge-0/0/7.0 and ge-0/0/8.0	Core interfaces on PE-2 that receive MPLS packets from the provider switch. The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.

[Table 32 on page 1332](#) shows the MPLS configuration components used for the provider switch in this example.

Table 32: CoS Configuration Components of the Provider Switch

Property	Settings	Description
Provider switch hardware	EX Series switch	Transit switch within the MPLS network configuration.
Custom EXP classifier	exp1	Name of the custom EXP classifier.
Custom EXP rewrite rule	e1	Name of the custom EXP rewrite rule.
Core interfaces receiving packets from other MPLS switches.	ge-0/0/5.0 and ge-0/0/6.0	Interfaces that connect the provider switch to the ingress PE switch (PE-1). The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.

Table 32: CoS Configuration Components of the Provider Switch (*continued*)

Property	Settings	Description
Core interfaces transmitting packets to other switches within the MPLS network.	ge-0/0/7.0 and ge-0/0/8.0	Interfaces that transmit packets to the egress PE (PE-2). The EXP rewrite rule is applied implicitly on these interfaces. Schedulers can also be specified and will be applied to these interfaces.

Configuring the Local PE Switch

CLI Quick Configuration

To quickly configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the local PE switch, copy the following commands and paste them into the switch terminal window of PE-1:

```
[edit]
set class-of-service classifiers dscp dscp1 import default
set class-of-service classifiers dscp dscp1 forwarding-class expedited-forwarding loss-priority low
code-points 000111
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
set firewall policer mypolicer if-exceeding bandwidth-limit 500m
set firewall policer mypolicer if-exceeding burst-size-limit 33553920
set firewall policer mypolicer then discard
set firewall family any filter myfilter term t1 then policer mypolicer
set protocols mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3 policing filter myfilter
```

Step-by-Step Procedure

To configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the ingress PE switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers dscp dscp1 import default
```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
```



```
user@switch# set classifiers dscp dscp1 forwarding-class expedited-forwarding loss-priority low
code-points 000111
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

4. Bind the DSCP classifier to the CCC interface:

```
[edit ]
user@switch# set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
```

5. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the LSP:

```
[edit firewall]
set policer mypolicer if-exceeding bandwidth-limit 500m
```

6. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

```
[edit firewall policer]
set mypolicer if-exceeding burst-size-limit 33553920
```

7. Discard traffic that exceeds the rate limits for this policer:

```
[edit firewall policer]
set mypolicer then discard
```

8. To reference the policer, configure a filter term that includes the policer action:

```
[edit firewall]
user@switch# set family any filter myfilter term t1 then policer mypolicer
```

9. Apply the filter to the LSP:

```
[edit protocols mpls]
set label-switched-path lsp_to_pe2_ge1 policing filter myfilter
```

Results

Display the results of the configuration:


```

[edit]
user@switch# show
class-of-service {
  classifiers {
    dscp dscp1 {
      import default;
      forwarding-class expedited-forwarding {
        loss-priority low code-points 000111;
      }
    }
  }
}
interfaces {
  ge-0/0/1 {
    unit 0 {
      classifiers {
        dscp dscp1;
      }
    }
  }
}
rewrite-rules {
  exp e1 {
    forwarding-class expedited-forwarding {
      loss-priority low code-point 111;
    }
  }
}
firewall {
  family any {
    filter myfilter {
      term t1 {
        then policer mypolicer;
      }
    }
  }
  policer mypolicer {
    if-exceeding {
      bandwidth-limit 500m;
      burst-size-limit 33553920;
    }
    then discard;
  }
}

```


Configuring the Remote PE Switch

CLI Quick Configuration

To quickly configure a custom EXP classifier on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE-2:

```
[edit]
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

Step-by-Step Procedure

To configure a custom EXP classifier on the egress PE switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

Results

Display the results of the configuration:

```
[edit]
user@switch# show
class-of-service {
  classifiers {
    exp exp1 {
      import default;
      forwarding-class expedited-forwarding {
        loss-priority low code-points 010;
      }
    }
  }
}
```


Configuring the Provider Switch

CLI Quick Configuration

To quickly configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch, copy the following commands and paste them into the switch terminal window of the provider switch:

```
[edit]
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

Step-by-Step Procedure

To configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

Results

Display the results of the configuration:

```
[edit]
user@switch# show
class-of-service {
  classifiers {
    exp exp1 {
```



```

import default;
forwarding-class expedited-forwarding {
    loss-priority low code-points 010;
}
}
rewrite-rules {
    exp e1 {
        forwarding-class expedited-forwarding {
            loss-priority low code-point 111;
        }
    }
}
}

```

Verification

IN THIS SECTION

- [Verifying That the Policer Firewall Filter Is Operational | 1338](#)
- [Verifying That the CoS Classifiers Are Going to the Right Queue | 1339](#)
- [Verifying the CoS Forwarding Table Mapping | 1343](#)
- [Verifying the Rewrite Rules | 1343](#)

To confirm that the configuration is working properly, perform these tasks:

Verifying That the Policer Firewall Filter Is Operational

Purpose

Verify the operational state of the policer that is configured on the ingress PE switch.

Action

user@switch> **show firewall**

```

Filter: myfilter
Policers:
Name                               Packets
mypolicer-t1                       0

```


Meaning

This output shows that the firewall filter **mypolicer** has been created.

Verifying That the CoS Classifiers Are Going to the Right Queue**Purpose**

Verify that the CoS classifiers are going to the right queue.

Action

```
user@switch> show class-of-service forwarding-table classifier
```

```
Classifier table index: 7, # entries: 64, Table type: DSCP
```

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	0	0
11	001011	0	0
12	001100	0	0
13	001101	0	0
14	001110	0	0
15	001111	0	0
16	010000	0	0
17	010001	0	0
18	010010	0	0
19	010011	0	0
20	010100	0	0
21	010101	0	0
22	010110	0	0
23	010111	0	0
24	011000	0	0
25	011001	0	0
26	011010	0	0
27	011011	0	0
28	011100	0	0
29	011101	0	0

30	011110	0	0
31	011111	0	0
32	100000	0	0
33	100001	0	0
34	100010	0	0
35	100011	0	0
36	100100	0	0
37	100101	0	0
38	100110	0	0
39	100111	0	0
40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0
45	101101	0	0
46	101110	0	0
47	101111	0	0
48	110000	3	0
49	110001	3	0
50	110010	3	0
51	110011	3	0
52	110100	3	0
53	110101	3	0
54	110110	3	0
55	110111	3	0
56	111000	3	0
57	111001	3	0
58	111010	3	0
59	111011	3	0
60	111100	3	0
61	111101	3	0
62	111110	3	0
63	111111	3	0

Classifier table index: 11, # entries: 8, Table type: IEEE 802.1

Entry #	Code point	Forwarding-class #	PLP
0	000	0	0
1	001	0	0
2	010	0	0
3	011	0	0
4	100	0	0
5	101	0	0
6	110	3	0


```

7          111          3          0

```

Classifier table index: 12, # entries: 8, Table type: IPv4 precedence

Entry #	Code point	Forwarding-class #	PLP
0	000	0	0
1	001	0	0
2	010	0	0
3	011	0	0
4	100	0	0
5	101	0	0
6	110	3	0
7	111	3	0

Classifier table index: 16, # entries: 8, Table type: Untrust

Entry #	Code point	Forwarding-class #	PLP
0	000	0	0
1	001	0	0
2	010	0	0
3	011	0	0
4	100	0	0
5	101	0	0
6	110	0	0
7	111	0	0

Classifier table index: 9346, # entries: 64, Table type: DSCP

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	1	0
8	001000	0	0
9	001001	0	0
10	001010	0	0
11	001011	0	0
12	001100	0	0
13	001101	0	0
14	001110	0	0
15	001111	0	0
16	010000	0	0
17	010001	0	0

18	010010	0	0
19	010011	0	0
20	010100	0	0
21	010101	0	0
22	010110	0	0
23	010111	0	0
24	011000	0	0
25	011001	0	0
26	011010	0	0
27	011011	0	0
28	011100	0	0
29	011101	0	0
30	011110	0	0
31	011111	0	0
32	100000	0	0
33	100001	0	0
34	100010	0	0
35	100011	0	0
36	100100	0	0
37	100101	0	0
38	100110	0	0
39	100111	0	0
40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0
45	101101	0	0
46	101110	0	0
47	101111	0	0
48	110000	3	0
49	110001	3	0
50	110010	3	0
51	110011	3	0
52	110100	3	0
53	110101	3	0
54	110110	3	0
55	110111	3	0
56	111000	3	0
57	111001	3	0
58	111010	3	0
59	111011	3	0
60	111100	3	0
61	111101	3	0

62	111110	3	0
63	111111	3	0

Meaning

This output shows that a new DSCP classifier has been created, index **9346**, on the ingress PE switch (PE-1).

Verifying the CoS Forwarding Table Mapping

Purpose

For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.

Action

```
user@switch>show class-of-service forwarding-table classifier mapping
```

Interface	Index	Table Index/	
		Q num	Table type
ge-0/0/1.0	92	9346	DSCP

Meaning

The results show that the new DSCP classifier, index number **9346**, is bound to interface **ge-0/0/1.0**.

Verifying the Rewrite Rules

Purpose

Display mapping of the queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.

Action

```
user@switch>show class-of-service forwarding-table rewrite-rule
```

```
Rewrite table index: 31, # entries: 4, Table type: DSCP
```

FC#	Low bits	State	High bits	State
0	000000	Enabled	000000	Enabled
1	101110	Enabled	101110	Enabled
2	001010	Enabled	001100	Enabled


```

3          110000  Enabled          111000  Enabled

Rewrite table index: 34, # entries: 4, Table type: IEEE 802.1
FC#      Low bits  State      High bits  State
0         000  Enabled          001  Enabled
1         010  Enabled          011  Enabled
2         100  Enabled          101  Enabled
3         110  Enabled          111  Enabled

Rewrite table index: 35, # entries: 4, Table type: IPv4 precedence
FC#      Low bits  State      High bits  State
0         000  Enabled          000  Enabled
1         101  Enabled          101  Enabled
2         001  Enabled          001  Enabled
3         110  Enabled          111  Enabled

Rewrite table index: 9281, # entries: 1, Table type: EXP
FC#      Low bits  State      High bits  State
1         111  Enabled          000  Disabled

```

Meaning

This output shows that a new EXP classifier with the index number **9281** has been created.

RELATED DOCUMENTATION

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

[Understanding Using CoS with MPLS Networks on EX Series Switches | 1324](#)

Monitoring CoS Forwarding Classes

Understanding CoS MPLS EXP Classifiers and Rewrite Rules

IN THIS SECTION

- [EXP Classifiers | 1346](#)
- [EXP Rewrite Rules | 1348](#)
- [Schedulers | 1348](#)

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, QFX3600, 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 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*

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 | 1345](#)

[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)

[Monitoring CoS Rewrite Rules](#)

[Defining CoS Rewrite Rules](#)

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 User Guide (Routers and EX9200 Switches)* 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 | 1345](#)

[Configuring a Global MPLS EXP Classifier | 1352](#)

[Configuring Rewrite Rules for MPLS EXP Classifiers | 1349](#)

Defining CoS Rewrite Rules

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 | 1345](#)

[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)

[Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)](#)

[Configuring Rewrite Rules for MPLS EXP Classifiers | 1349](#)

Configuring Generalized MPLS (GMPLS)

IN THIS CHAPTER

- [Introduction to GMPLS | 1353](#)
- [GMPLS Terms and Acronyms | 1355](#)
- [Supported GMPLS Standards | 1356](#)
- [GMPLS Operation | 1357](#)
- [GMPLS and OSPF | 1358](#)
- [GMPLS and CSPF | 1358](#)
- [GMPLS Features | 1359](#)
- [LMP Configuration Overview | 1359](#)
- [Configuring LMP Traffic Engineering Links | 1361](#)
- [Configuring LMP Peers | 1363](#)
- [Configuring RSVP and OSPF for LMP Peer Interfaces | 1369](#)
- [Configuring MPLS Paths for GMPLS | 1371](#)
- [Tracing LMP Traffic | 1371](#)
- [Configuring MPLS LSPs for GMPLS | 1372](#)
- [Gracefully Tearing Down GMPLS LSPs | 1375](#)
- [GMPLS RSVP-TE VLAN LSP Signaling Overview | 1377](#)
- [Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling | 1384](#)

Introduction to GMPLS

Traditional MPLS is designed to carry Layer 3 IP traffic using established IP-based paths and associating these paths with arbitrarily assigned labels. These labels can be configured explicitly by a network administrator, or can be dynamically assigned by means of a protocol such as LDP or RSVP.

GMPLS generalizes MPLS in that it defines labels for switching varying types of Layer 1, Layer 2, or Layer 3 traffic. GMPLS nodes can have links with one or more of the following switching capabilities:

- Fiber-switched capable (FSC)
- Lambda-switched capable (LSC)
- Time-division multiplexing (TDM) switched-capable (TSC)
- Packet-switched capable (PSC)

Label-switched paths (LSPs) must start and end on links with the same switching capability. For example, routers can establish packet-switched LSPs with other routers. The LSPs might be carried over a TDM-switched LSP between SONET add/drop multiplexers (ADMs), which in turn might be carried over a lambda-switched LSP.

The result of this extension of the MPLS protocol is an expansion in the number of devices that can participate in label switching. Lower-layer devices, such as OXCs and SONET ADMs, can now participate in GMPLS signaling and set up paths to transfer data. A router can participate in signaling optical paths across a transport network.

Two service models determine the visibility that a client node (a router, for example) has into the optical core or transport network. The first is through a user-to-network interface (UNI), which is often referred to as the overlay model. The second is known as the peer model. Juniper Networks supports both models.

NOTE: There is not necessarily a one-to-one correspondence between a physical interface and a GMPLS interface. If a GMPLS connection uses a nonchannelized physical connector, the GMPLS label can use the physical port ID. However, the label for channelized interfaces often is based on a channel or time slot. Consequently, it is best to refer to GMPLS labels as identifiers for a resource on a traffic engineering link.

To establish LSPs, GMPLS uses the following mechanisms:

- An out-of-band control channel and a data channel—RSVP messages for LSP setup are sent over an out-of-band control network. Once the LSP setup is complete and the path is provisioned, the data channel is up and can be used to carry traffic. The Link Management Protocol (LMP) is used to define and manage the data channels between a pair of nodes. You can optionally use LMP to establish and maintain LMP control channels between peers running the same Junos OS Release.
- RSVP-TE extensions for GMPLS—RSVP-TE is already designed to signal the setup of packet LSPs. This has been extended for GMPLS to be able to request path setup for various kinds of LSPs (nonpacket) and request labels like wavelengths, time slots, and fibers as label objects.
- Bidirectional LSPs—Data can travel both ways between GMPLS devices over a single path, so nonpacket LSPs are signaled to be bidirectional.

GMPLS Terms and Acronyms

F

Forwarding adjacency A forwarding path for sending data between GMPLS-enabled devices.

G

Generalized MPLS (GMPLS) An extension to MPLS that allows data from multiple layers to be switched over label-switched paths (LSPs). GMPLS LSP connections are possible between similar Layer 1, Layer 2, and Layer 3 devices.

GMPLS label Layer 3 identifiers, fiber port, time-division multiplexing (TDM) time slot, or dense wavelength-division multiplexing (DWDM) wavelength of a GMPLS-enabled device used as a next-hop identifier.

GMPLS LSP types The four types of GMPLS LSPs are:

- Fiber-switched capable (FSC)—LSPs are switched between two fiber-based devices, such optical cross-connects (OXC) that operate at the level of individual fibers.
- Lambda-switched capable (LSC)—LSPs are switched between two DWDM devices, such as such as OXC that operate at the level of individual wavelengths.
- TDM-switched capable (TDM)—LSPs are switched between two TDM devices, such as SONET ADMs.
- Packet-switched capable (PSC)—LSPs are switched between two packet-based devices, such as routers or ATM switches.

L

Link Management Protocol A protocol used to define a forwarding adjacency between peers and to maintain and allocate resources on the traffic engineering links.

T

Traffic engineering link A logical connection between GMPLS-enabled devices. Traffic engineering links can have addresses or IDs and are associated with certain resources or interfaces. They also have certain attributes (encoding-type, switching capability, bandwidth, and so on). The logical addresses can be routable, although this is not required because they are acting as link identifiers. Each traffic engineering link represents a forwarding adjacency between a pair of devices.

Supported GMPLS Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for Generalized MPLS (GMPLS).

- RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

Only the following features are supported:

- Bidirectional LSPs (upstream label only)
- Control channel separation
- Generalized label (suggested label only)
- Generalized label request (bandwidth encoding only)

- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

Only Section 9, "Fault Handling," is supported.

- RFC 4202, *Routing Extensions in Support of Generalized Multi-Protocol Label Switching*

Only interface switching is supported.

- RFC 4206, *Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*

- Internet draft draft-ietf-ccamp-gmpls-rsvp-te-ason-02.txt, *Generalized MPLS (GMPLS) RSVP-TE Signalling in support of Automatically Switched Optical Network (ASON)* (expires January 2005)

- Internet draft draft-ietf-ccamp-gmpls-sonet-sdh-08.txt, *Generalized Multi-Protocol Label Switching Extensions for SONET and SDH Control*

Only S,U,K,L,M-format labels and SONET traffic parameters are supported.

- Internet draft draft-ietf-ccamp-lmp-10.txt, *Link Management Protocol (LMP)*

- Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-12.txt, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching*

The following sub-TLV types for the Link type, link, value (TLV) are not supported:

- Link Local/Remote Identifiers (type 11)
- Link Protection Type (type 14)
- Shared Risk Link Group (SRLG) (type 16)

The features described in Section 2 of the draft, “Implications on Graceful Restart,” are also not supported.

The Interface Switching Capability Descriptor (type 15) sub-TLV type is implemented, but only for packet switching.

- Internet draft draft-ietf-mpls-bundle-04.txt, *Link Bundling in MPLS Traffic Engineering*

RELATED DOCUMENTATION

[Supported LDP Standards | 923](#)

[Supported MPLS Standards | 20](#)

[Supported RSVP Standards | 816](#)

[Accessing Standards Documents on the Internet](#)

GMPLS Operation

The basic functionality of GMPLS requires close interaction between RSVP and LMP. It works in the following sequence:

1. LMP notifies RSVP of the new entities:
 - Traffic engineering link (forwarding adjacency)
 - Resources available for the traffic engineering link
 - Control peer
2. GMPLS extracts the LSP attributes from the configuration and requests RSVP to signal one or more specific paths, which are specified by the traffic engineering link addresses.
3. RSVP determines the local traffic engineering link, corresponding control adjacency and active control channel, and transmission parameters (such as IP destination). It requests that LMP allocate a resource from the traffic engineering link with the specified attributes. If LMP finds a resource matching the

attributes, label allocation succeeds. RSVP sends a PathMsg hop by hop until it reaches the target router.

4. When the target router receives the PathMsg, RSVP again requests that LMP allocate a resource based on the signaled parameters. If label allocation succeeds, the router sends back a ResvMsg.
5. If the signaling is successful, a bidirectional optical path is provisioned.

GMPLS and OSPF

You can configure OSPF for GMPLS. OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions.

GMPLS and CSPF

GMPLS introduces extra constraints for computing paths for GMPLS LSPs that use CSPF. These additional constraints affect the following link attributes:

- Signal type (minimum LSP bandwidth)
- Encoding type
- Switching type

These new constraints are populated in the traffic engineering database with the exchange of an interface-switching capability descriptor type, length, value (TLV) through an IGP.

The ignored constraints that are exchanged through the interface switching capability descriptor include:

- Maximum LSP bandwidth
- Maximum transmission unit (MTU)

The CSPF path computation is the same as in non-GMPLS environments, except that the links are also limited by GMPLS constraints.

Each link can have multiple interface-switching capability descriptors. All the descriptors are checked before a link is rejected.

The constraints are checked in the following order:

1. The signal type configured for the GMPLS LSP signifies the amount of bandwidth requested. If the desired bandwidth is less than the minimum LSP bandwidth, the interface-switching descriptor is rejected.

2. The encoding type of the link for the ingress and the egress interfaces should match. The encoding type is selected and stored at the ingress node after all the constraints are satisfied by the link and is used to select the link on the egress node.
3. The switching type of the links of the intermediate switches should match that of the GMPLS LSP specified in the configuration.

GMPLS Features

The Junos OS includes the following GMPLS functionality:

- An out-of-band control plane makes it possible to signal LSP path setup.
- RSVP-TE extensions support additional objects beyond Layer 3 packets, such as ports, time slots, and wavelengths.
- The LMP protocol creates and maintains a database of traffic engineering links and peer information. Only the static version of this protocol is supported in the Junos OS. You can optionally configure LMP to establish and maintain LMP control channels between peers running the same Junos OS Release.
- Bidirectional LSPs are required between devices.
- Several GMPLS label types that are defined in RFC 3471, *Generalized MPLS—Signaling Functional Description*, such as MPLS, Generalized, SONET/SDH, Suggested, and Upstream, are supported. Generalized labels do not contain a type field, because the nodes should know from the context of their connection what type of label to expect.
- Traffic parameters facilitate GMPLS bandwidth encoding and SONET/SDH formatting.
- Other supported attributes include interface identification and errored interface identification, user-to-network (UNI)-style signaling, and secondary LSP paths.

LMP Configuration Overview

You need to configure the Link Management Protocol (LMP) to define the data channel connection and the control channel connection between devices. Include the following statements at the **[edit protocols link-management]** hierarchy level:

```
[edit protocols link-management]
peer peer-name {
  address address;
  control-channel control-channel-name;
```



```

lmp-control-channel control-channel-interface {
    remote-address ip-address;
}
lmp-protocol {
    hello-dead-interval milliseconds;
    hello-interval milliseconds;
    retransmission-interval milliseconds;
    retry-limit number;
    passive;
}
te-link te-link-name;
}
te-link te-link-name {
    disable;
    interface interface-name {
        disable;
        local-address ip-address;
        remote-address ip-address;
        remote-id id-number;
    }
    label-switched-path lsp-name;
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}

```

NOTE: Although you can include GMPLS configuration statements at the **[edit logical-systems logical-system-name]** hierarchy level, GMPLS is not supported on logical systems.

For information about configuring LMP, see the following sections:

- [Configuring LMP Traffic Engineering Links on page 1361](#)
- [Configuring LMP Peers on page 1363](#)
- [Configuring RSVP and OSPF for LMP Peer Interfaces on page 1369](#)
- [Configuring MPLS Paths for GMPLS on page 1371](#)
- [Tracing LMP Traffic on page 1371](#)

Configuring LMP Traffic Engineering Links

IN THIS SECTION

- [Configuring the Local IP Address for Traffic Engineering Links | 1361](#)
- [Configuring the Remote IP Address for Traffic Engineering Links | 1362](#)
- [Configuring the Remote ID for Traffic Engineering Links | 1362](#)

An LMP traffic engineering link acts as a data channel connection between GMPLS devices.

To configure a traffic engineering link, include the **te-link** statement at the **[edit protocols link-management]** hierarchy level:

```
[edit protocols link-management]
te-link te-link-name {
  disable;
  interface interface-name {
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
  }
  label-switched-path lsp-name;
  local-address ip-address;
  remote-address ip-address;
  remote-id id-number;
}
```

Complete the procedures in the following sections to configure an LMP traffic engineering link:

When you configure a traffic engineering link that contains interfaces for an LMP peer, you must also configure a control channel. However, no control channel is required for a traffic engineering link that contains an LSP. For information about configuring control channels, see [“Configuring LMP Peers” on page 1363](#).

Configuring the Local IP Address for Traffic Engineering Links

Use the **local-address** statement to configure the local IP address associated with the traffic engineering link.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This configuration enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the local IP address for the traffic engineering link, include the **local-address** statement:

```
te-link te-link-name {
  interface interface-name {
    local-address ip-address;
  }
  local-address ip-address;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Remote IP Address for Traffic Engineering Links

You need to specify the address of the remote end of the data channel for each traffic engineering link. Use the **remote-address** statement to configure the remote IP address.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the remote IP address for the traffic engineering link, include the **remote-address** statement:

```
te-link te-link-name {
  interface interface-name {
    remote-address ip-address;
  }
  remote-address ip-address;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring the Remote ID for Traffic Engineering Links

The local ID for the traffic engineering link is automatically assigned by LMP. The port identifier and labels for the interfaces (resources) in the traffic engineering link are also assigned automatically. However, you need to explicitly configure the remote ID for the traffic engineering link and the remote ID traffic engineering link interface. The remote ID for the interface must be based on the post-ID assignment of the peer node. The remote IDs are needed for static mapping of remote labels to local labels.

Before you can obtain the remote IDs for the traffic engineering link and traffic engineering link interface on the peer node, you must first configure the LMP peer, as described in [“Configuring LMP Peers” on page 1363](#). Once you have configured the LMP peer, you can obtain the traffic engineering link local ID and interface local ID by issuing the **show link-management te-link** command. Once you have these IDs, you can configure them as the remote IDs on the peer node.

To configure the remote ID for a traffic engineering link and for the traffic engineering link interface, include the **remote-id** statement:

```
te-link te-link-name {
  interface interface-name {
    remote-id id-number;
  }
  remote-id id-number;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Configuring LMP Peers

IN THIS SECTION

- [Configuring the ID for LMP Peers | 1364](#)
- [Configuring the Interface for Control Channels Between LMP Peers | 1364](#)
- [Configuring the LMP Control Channel Interface for the Peer | 1365](#)
- [Configuring the Remote IP Address for LMP Control Channels | 1366](#)
- [Configuring Hello Message Intervals for LMP Control Channels | 1366](#)
- [Controlling Message Exchange for LMP Control Channels | 1367](#)
- [Preventing the Local Peer from Initiating LMP Negotiation | 1367](#)
- [Associating Traffic Engineering Links with LMP Peers | 1368](#)
- [Disabling the Traffic Engineering Link for LMP Peers | 1368](#)

You need to configure network peers for GMPLS. A peer is a network device that your router communicates with when setting up the control and data channels. The peer is often an optical cross-connect (OXC).

To configure an LMP peer name, include the **peer** statement at the [edit protocols link-management] hierarchy level:

```
peer peer-name {
  address ip-address;
  control-channel control-channel-interface;
  lmp-control-channel control-channel-interface {
    remote-address ip-address;
  }
  lmp-protocol {
    hello-dead-interval milliseconds;
    hello-interval milliseconds;
    retransmission-interval milliseconds;
    retry-limit number;
    passive;
  }
  te-link te-link-name;
}
```

The following sections describe how to configure an LMP peer:

Configuring the ID for LMP Peers

To configure the LMP peer ID, include the **address** statement at the [edit protocols link-management peer *peer-name*] hierarchy level. The default value for the LMP peer ID is the loopback address.

```
[edit protocols link-management peer peer-name]
address ip-address;
```

Configuring the Interface for Control Channels Between LMP Peers

You must configure one or more control channels between the LMP peers. The control channels must travel across either a point-to-point link or a tunnel.

To configure the interface for the control channel, include the **control-channel** statement at the [edit protocols link-management peer *peer-name*] hierarchy level:

```
[edit protocols link-management peer peer-name]
control-channel [ interface-names ];
```

You can configure a generic routing encapsulation (GRE) interface (gre-x/y/z) for the control channel. This type of interface does not require a Tunnel PIC.

NOTE: You can configure GRE interfaces (gre-x/y/z) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

Configuring the LMP Control Channel Interface for the Peer

In an environment that uses LMP to establish and maintain an LMP control channel between peers, you can configure a number of attributes associated with LMP. To configure the interface to be associated with the LMP control channel for the peer, include the **lmp-control-channel** statement:

```
lmp-control-channel control-channel-interface;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name*]
- [edit logical-systems *logical-system-name* protocols link-management **peer** *peer-name*]

You can configure a GRE interface for the LMP control channel. This type of interface does not require a Tunnel PIC.

NOTE: You can configure GRE interfaces only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

When this LMP control channel interface comes up, the peers use LMP to negotiate channel parameters and configure the control channel.

The local peer repeatedly sends a Config message to the remote peer. The Config message contains the local control channel ID, the local peer's node ID, a message ID, and a CONFIG object that includes hello message attributes (the hello interval and the hello dead interval).

The channel is activated when the remote peer responds with a ConfigAck message. The remote peer does so only when its own configured hello interval and hello dead interval match the values in the received Config message or the default values. If these values do not match, the remote peer responds with a ConfigNack message. The local peer logs this event and resends the Config message until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the configuration process.

Configuring the Remote IP Address for LMP Control Channels

You need to specify the address of the remote end of the LMP control channel.

To configure the remote IP address for the LMP control channel, include the **remote-address** statement:

```
remote-address address;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name* **lmp-control-channel** *control-channel-interface*]
- [edit logical-systems *logical-system-name* protocols link-management **peer** *peer-name* **lmp-control-channel** *control-channel-interface*]

Configuring Hello Message Intervals for LMP Control Channels

Hello messages are exchanged between LMP peers to maintain the control channel after LMP has activated the control channel. The LMP control channel is considered to be up only when the hello negotiation is successful. Successful negotiation consists of the local peer sending a hello message to the remote peer and receiving a hello message in response.

The LMP peers continue to exchange hello messages after the LMP control channel is up in order to maintain the channel.

The hello interval specifies the interval between periodic hello messages. The hello dead interval specifies how long the local peer waits for a hello response before it declares the LMP control channel to be down. When the channel goes down, the local peer restarts the LMP control channel negotiation and configuration process.

You can specify a hello interval from 150 through 300,000 milliseconds. The default hello interval is 150 milliseconds.

You can specify a hello dead interval from 500 through 300,000 milliseconds. The default hello dead interval is 500 milliseconds.

To configure the attributes for hello messages exchanged between LMP peers, include the **hello-interval** and **hello-dead-interval** statements:

```
hello-dead-interval milliseconds;  
hello-interval milliseconds;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name* **lmp-protocol**]

- [edit logical-systems *logical-system-name* protocols link-management *peer peer-name* *lmp-protocol*]

When an LMP control channel comes up after a successful exchange of hello messages between LMP peers, LMP uses link property correlation to verify the traffic engineering and data link information on both sides of a link. To do so, the local peer sends a LinkSummary message for each traffic engineering link governed by the LMP control channel. The LinkSummary message contains information that characterizes the traffic engineering link and each data link in the traffic engineering link.

The local peer continues sending a LinkSummary message for each link until the remote peer responds with a LinkSummaryAck message or until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the link property correlation process.

When the remote peer receives a LinkSummary message, it examines its own link information. If this information agrees with that in the LinkSummary message, the remote peer responds with a LinkSummaryAck message. If the information is different, the remote peer responds with a LinkSummaryNack message.

Controlling Message Exchange for LMP Control Channels

You can configure message attributes that control the exchange of LMP Config and LinkSummary messages. The retransmission interval specifies the interval between resubmitted LMP messages. The retry limit specifies how many times LMP sends a message before restarting the process.

You can specify a retransmission interval from 500 through 300,000 milliseconds. The default retransmission interval is 500 milliseconds.

You can specify a retry limit from 3 through 1000 attempts. The default number of retry attempts is three.

To configure attributes governing the exchange of LMP messages between peers, include the **retransmission-interval** and **retry-limit** statements:

```
retransmission-interval milliseconds;  
retry-limit number;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management *peer peer-name* *lmp-protocol*]
- [edit logical-systems *logical-system-name* protocols link-management *peer peer-name* *lmp-protocol*]

Preventing the Local Peer from Initiating LMP Negotiation

You can specify that the local peer does not initiate LMP negotiation. Instead, the local peer waits for the remote peer to configure the LMP control channel.

To configure the local peer to wait for the remote peer to configure the LMP control channel, include the **passive** statement:

```
passive;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name* **lmp-protocol**]
- [edit logical-systems *logical-system-name* protocols link-management **peer** *peer-name* **lmp-protocol**]

Associating Traffic Engineering Links with LMP Peers

To specify the name of a traffic engineering link to be associated with this peer, include the **te-link** statement at the [edit protocols link-management **peer** *peer-name*] hierarchy level:

```
[edit protocols link-management peer peer-name]  
te-link te-link-name;
```

For information about how to configure a traffic engineering link, see “[Configuring LMP Traffic Engineering Links](#)” on page 1361.

Disabling the Traffic Engineering Link for LMP Peers

To disable a specific traffic engineering link, include the **disable** statement:

```
disable;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **te-link** *te-link-name*]
- [edit logical-systems *logical-system-name* protocols link-management **te-link** *te-link-name*]

Configuring RSVP and OSPF for LMP Peer Interfaces

IN THIS SECTION

- [Configuring RSVP Signaling for LMP Peer Interfaces | 1369](#)
- [Configuring OSPF Routing for LMP Peer Interfaces | 1370](#)
- [Configuring the Hello Interval for LMP Peer Interfaces | 1370](#)

After you have configured the LMP peers as described in [“Configuring LMP Peers” on page 1363](#), add the peer interfaces to RSVP and OSPF. The peer interface name must match the peer name configured in LMP. Once the peer interfaces are added to the protocols, the traffic engineering link local and remote addresses can be signaled and advertised to peers like any other interface enabled for RSVP and OSPF. These addresses act as virtual interfaces for GMPLS.

NOTE: When adding the virtual peer interfaces to RSVP and OSPF, do not configure the corresponding physical control channel interface in either protocol. If you include the **interface all** statement, you must disable RSVP and OSPF protocols manually on the control channel interface.

To configure peer interfaces in RSVP and OSPF, complete the procedures in the following sections:

Configuring RSVP Signaling for LMP Peer Interfaces

To configure RSVP signaling for LMP peers, configure the LMP peer interface by including the **peer-interface** statement at the **[edit protocols rsvp]** hierarchy level:

```
[edit protocols rsvp]
peer-interface peer-interface-name {
  (aggregate | no-aggregate);
  authentication-key key;
  disable;
  hello-interval seconds;
  (reliable | no-reliable);
}
```


The statements configured at the `[edit protocols rsvp peer-interface peer-interface-name]` hierarchy level have the same functionality as the statements configured at the `[edit protocols rsvp interface interface-name]` hierarchy level.

Configuring OSPF Routing for LMP Peer Interfaces

To configure OSPF routing for LMP peers, configure the name of the LMP peer by including the `peer-interface` statement at the `[edit protocols ospf area area-number]` hierarchy level:

```
[edit protocols ospf area area-number]
peer-interface peer-interface-name {
  dead-interval seconds;
  disable;
  hello-interval seconds;
  retransmit-interval seconds;
  transit-delay seconds;
}
```

For information about how to configure OSPF statements, see the *Junos OS Routing Protocols Library*.

Configuring the Hello Interval for LMP Peer Interfaces

Hello packets are used to indicate to neighboring routers that the peer interface is still up and running. The hello interval must be the same for all routers on a shared logical IP network. You can specify a hello interval from 1 through 255 seconds. The default hello interval is normally 10 seconds. For nonbroadcast networks, the default hello interval is 120 seconds.

To specify how often the router sends hello packets out the peer interface, configure the `hello-interval` statement:

```
hello-interval seconds;
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols ospf area area-number peer-interface peer-interface-name]`
- `[edit logical-systems logical-system-name protocols ospf area area-number peer-interface peer-interface-name]`

Configuring MPLS Paths for GMPLS

As part of the configuration for GMPLS, you need to establish an MPLS path for each unique device connected through GMPLS. Configure the traffic engineering link remote address as the address at the `[edit protocols mpls path path-name]` hierarchy level. Constrained Shortest Path First (CSPF) is supported so you can choose either the **strict** or **loose** option with the address.

See [“LMP Configuration Overview” on page 1359](#) for information about how to obtain a traffic engineering link remote address.

To configure the MPLS path, include the **path** statement at the `[edit protocols mpls]` hierarchy level:

```
[edit protocols mpls]
path path-name {
    next-hop-address (strict | loose);
}
```

For information about how to configure MPLS paths, see [“Creating Named Paths” on page 516](#).

Tracing LMP Traffic

To trace LMP protocol traffic, include the **traceoptions** statement at the `[edit protocols link-management]` hierarchy level:

```
[edit protocols link-management]
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`.

The following trace flags display the operations associated with the sending and receiving of various LMP messages:

- **all**—Trace all available operations
- **hello-packets**—Trace hello packets on any LMP control channel
- **init**—Output from the initialization messages
- **packets**—Trace all packets other than hello packets on any LMP control channel

- **parse**—Operation of the parser
- **process**—Operation of the general configuration
- **route-socket**—Operation of route socket events
- **routing**—Operation of the routing protocols
- **server**—Server processing operations
- **show**—Servicing operations for **show** commands
- **state**—Trace state transitions of the LMP control channels and traffic engineering links

Each flag can carry one or more of the following flag modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

Configuring MPLS LSPs for GMPLS

IN THIS SECTION

- [Configuring the Encoding Type | 1373](#)
- [Configuring the GPID | 1373](#)
- [Configuring the Signal Bandwidth Type | 1374](#)
- [Configuring GMPLS Bidirectional LSPs | 1374](#)
- [Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS | 1374](#)

To enable the proper GMPLS switching parameters, configure the label-switched path (LSP) attributes that are appropriate for your network connection. The default value for **switching-type** is **psc-1**, which is also appropriate for standard MPLS.

To configure the LSP attributes, include the **lsp-attributes** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name]
lsp-attributes {
  encoding-type type;
```



```

gpid gpid;
signal-bandwidth type;
switching-type type;
}

```

If you include the **no-cspf** statement in the label-switched path configuration, you must also configure primary and secondary paths, or the configuration cannot be committed.

The following sections describe how to configure each of the LSP attributes for a GMPLS LSP:

Configuring the Encoding Type

You need to specify the encoding type of the payload carried by the LSP. It can be any of the following:

- **ethernet**—Ethernet
- **packet**—Packet
- **pdh**—Plesiochronous digital hierarchy (PDH)
- **sonet-sdh**—SONET/SDH

The default value is **packet**.

To configure the encoding type, include the **encoding-type** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```

[edit protocols mpls label-switched-path lsp-name lsp-attributes]
encoding-type type;

```

Configuring the GPID

You need to specify the type of payload carried by the LSP. The payload is the type of packet underneath the MPLS label. The payload is specified by the generalized payload identifier (GPID).

You can specify the GPID with any of the following values:

- **hdlc**—High-Level Data Link Control (HDLC)
- **ethernet**—Ethernet
- **ipv4**—IP version 4 (default)
- **pos-scrambling-crc-16**—For interoperability with other vendors' equipment
- **pos-no-scrambling-crc-16**—For interoperability with other vendors' equipment
- **pos-scrambling-crc-32**—For interoperability with other vendors' equipment

- **pos-no-scrambling-crc-32**—For interoperability with other vendors' equipment
- **ppp**—Point-to-Point Protocol (PPP)

To configure the GPID, include the **gpipid** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
  gpipid gpipid;
```

Configuring the Signal Bandwidth Type

The signal bandwidth type is the encoding used for path computation and admission control. To configure the signal bandwidth type, include the **signal-bandwidth** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
  signal-bandwidth type;
```

Configuring GMPLS Bidirectional LSPs

Because MPLS and GMPLS use the same configuration hierarchy for LSPs, it is helpful to know which LSP attributes control LSP functionality. Standard MPLS packet-switched LSPs are unidirectional, whereas GMPLS nonpacket LSPs are bidirectional.

If you use the default packet-switching type of **psc-1**, your LSP becomes unidirectional. To enable a GMPLS bidirectional LSP, you must select a non-packet-switching type option, such as **lambda**, **fiber**, or **ethernet**. Include the **switching-type** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
  switching-type (lambda | fiber | ethernet);
```

Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS

By setting the A-bit in the Admin Status object, you can enable nonpacket GMPLS LSPs to establish paths through routers that run Junos. When an ingress router sends an RSVP PATH message with the Admin Status A-bit set, an external device (not a router running the Junos OS) can either perform a Layer 1 path setup test or help bring up an optical cross-connect.

When set, the A-bit in the Admin Status object indicates the administrative down status for a GMPLS LSP. This feature is used specifically by nonpacket GMPLS LSPs. It does not affect control path setup or data forwarding for packet LSPs.

Junos does not distinguish between the control path setup and data path setup. Other nodes along the network path use RSVP PATH signaling using the A-bit in a meaningful way.

To configure the Admin Status object for a GMPLS LSP, include the **admin-down** statement:

```
admin-down;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]

Gracefully Tearing Down GMPLS LSPs

IN THIS SECTION

- [Temporarily Deleting GMPLS LSPs | 1375](#)
- [Permanently Deleting GMPLS LSPs | 1376](#)
- [Configuring the Graceful Deletion Timeout Interval | 1376](#)

You can gracefully tear down nonpacket GMPLS LSPs. An LSP that is torn down abruptly, a common process in a packet-switched network, can cause stability problems in nonpacket-switched networks. To maintain the stability of nonpacket-switched networks, it might be necessary to tear down LSPs gracefully.

The following sections describe how to tear down GMPLS LSPs gracefully:

Temporarily Deleting GMPLS LSPs

You can gracefully tear down a GMPLS LSP using the **clear rsvp session gracefully** command.

This command gracefully tears down an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin Status object is signaled along the path to the endpoint of the LSP. During the second pass, the LSP is taken down. Using this command, the LSP is taken down temporarily. After the appropriate interval, the GMPLS LSP is resignaled and then reestablished.

The **clear rsvp session gracefully** command has the following properties:

- It only works on the ingress and egress routers of an RSVP session. If used on a transit router, it has the same behavior as the **clear rsvp session** command.
- It only works for nonpacket LSPs. If used with packet LSPs, it has the same behavior as the **clear rsvp session** command.

For more information, see the [CLI Explorer](#).

Permanently Deleting GMPLS LSPs

When you disable an LSP in the configuration, the LSP is permanently deleted. By configuring the **disable** statement, you can disable a GMPLS LSP permanently. If the LSP being disabled is a nonpacket LSP, then the graceful LSP tear-down procedures that use the Admin Status object are used. If the LSP being disabled is a packet LSP, then the regular signaling procedures for LSP deletion are used.

To disable a GMPLS LSP, include the **disable** statement at any of the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name*]**—Disable the LSP.
- **[edit protocols link-management *te-link* *te-link-name*]**—Disable a traffic engineering link.
- **[edit protocols link-management *te-link* *te-link-name* *interface* *interface-name*]**—Disable an interface used by a traffic engineering link.

Configuring the Graceful Deletion Timeout Interval

The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

The ingress router initiates the graceful deletion procedure by sending the Admin Status object in the path message with the **D** bit set. The ingress router expects to receive an Resv message with the **D** bit set from the egress router. If the ingress router does not receive this message within the time specified by the graceful deletion timeout interval, it initiates a forced tear-down of the LSP by sending a PathTear message.

To configure the graceful deletion timeout interval, include the **graceful-deletion-timeout** statement at the **[edit protocols rsvp]** hierarchy level. You can configure a time between 1 through 300 seconds. The default value is 30 seconds.

```
graceful-deletion-timeout seconds;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols rsvp]**

- [edit logical-systems *logical-system-name* protocols rsvp]

You can use the **show rsvp version** command to determine the current value configured for the graceful deletion timeout.

GMPLS RSVP-TE VLAN LSP Signaling Overview

IN THIS SECTION

- [Understanding GMPLS RSVP-TE Signaling | 1377](#)
- [Need for GMPLS RSVP-TE VLAN LSP Signaling | 1378](#)
- [GMPLS RSVP-TE VLAN LSP Signaling Functionality | 1379](#)
- [LSP Hierarchy with GMPLS RSVP-TE VLAN LSP | 1380](#)
- [Path Specification for GMPLS RSVP-TE VLAN LSP | 1380](#)
- [GMPLS RSVP-TE VLAN LSP Configuration | 1380](#)
- [Associated Bidirectional Packet LSP | 1382](#)
- [Make-Before-Break for Associated Bidirectional Packet and GMPLS RSVP-TE VLAN LSP | 1382](#)
- [Supported and Unsupported Features | 1383](#)

Understanding GMPLS RSVP-TE Signaling

Signaling is the process of exchanging messages within the control plane to set up, maintain, modify, and terminate data paths (label-switched paths (LSPs)) in the data plane. Generalized MPLS (GMPLS) is a protocol suite that extends the existing control plane of MPLS to manage further classes of interfaces and to support other forms of label switching, such as time-division multiplexing (TDM), fiber (port), Lambda, and so on.

GMPLS extends intelligent IP/MPLS connections from Layer 2 and Layer 3 all the way to Layer 1 optical devices. Unlike MPLS, which is supported mainly by routers and switches, GMPLS can also be supported by optical platforms, including SONET/SDH, optical cross-connects (OXCs), and dense wave division multiplexing (DWDM).

In addition to labels, which are primarily used to forward data in MPLS, other physical entries, such as wavelengths, time slots, and fibers can be used as label objects to forward data in GMPLS, thereby leveraging the existing control plane mechanisms to signal different kinds of LSPs. GMPLS uses RSVP-TE to be able to request the other label objects to signal the various kinds of LSPs (nonpacket). Bidirectional LSPs and

an out-of-band control channel and a data channel using the Link Management Protocol (LMP) are the other mechanisms that are used by GMPLS to establish LSPs.


Need for GMPLS RSVP-TE VLAN LSP Signaling

The traditional Layer 2 point-to-point services use Layer 2 circuits and Layer 2 VPN technologies that are based on LDP and BGP. In the traditional deployment, the customer edge (CE) devices do not participate in the signaling of the Layer 2 service. The provider edge (PE) devices manage and provision the Layer 2 service to provide end-to-end connectivity between the CE devices.

One of the biggest challenges of having the PE devices provision the Layer 2 services for each Layer 2 circuit between a pair of CE devices is the network management burden on the provider network.

Figure 105 on page 1378 illustrates how the Layer 2 service is set up and used by the CE routers in a LDP/BGP-based Layer 2 VPN technology. Two CE routers CE1 and CE2 are connected to a provider MPLS network through the PE routers PE1 and PE2 respectively. The CE routers are connected to the PE routers by Ethernet links. Routers CE1 and CE2 are configured with VLAN1 and VLAN2 logical Layer 3 interfaces, so they appear to be directly connected. Routers PE1 and PE2 are configured with Layer 2 circuit (pseudowire) to carry the Layer 2 VLAN traffic between the CE routers. The PE routers use packet MPLS LSPs within the provider MPLS network to carry the Layer 2 VLAN traffic.

Figure 105: Traditional Layer 2 Point-to-Point Services




With the introduction of GMPLS-based VLAN LSP signaling, the need for the PE (also called server-layer) network to provision each individual Layer 2 connection between the CE (also called client) devices is minimized. The client router requests the server-layer router to which it is directly connected, for setting up the Layer 2 service to connect with a remote client router through GMPLS signaling.

The server-layer devices extend the signaling through the server-layer network to connect with the remote client routers. In the process, the server-layer device sets up the data plane for the Layer 2 service at the server-client border, and sets up the data plane for carrying the Layer 2 traffic within the server-layer network. With the Layer 2 service setup, the client routers can run IP/MPLS directly on top of the Layer 2 service and have IP/MPLS adjacency with each other.

In addition to reducing the provisioning activity needed on the server-layer devices, GMPLS signaling also provides the client routers with the flexibility of bringing up the Layer 2 circuits on an on-demand basis without depending on the server-layer administration for the provisioning of the Layer 2 service.

Using the same topology as in Figure 1, [Figure 106 on page 1379](#) illustrates how the Layer 2 service is set up and used by the client routers in GMPLS RSVP-TE-based Layer 2 VPN technology.

Figure 106: GMPLS RSVP-TE VLAN LSP



In [Figure 106 on page 1379](#), instead of configuring a pseudowire to carry the Layer 2 VLAN traffic between the client routers, Routers PE1 and PE2 are configured with an IP-based communication channel and other GMPLS-specific configurations (identification of Ethernet links as TE-links) for allowing the exchange of GMPLS RSVP-TE signaling messages with the client routers. Routers CE1 and CE2 are also configured with an IP-based communication channel and relevant GMPLS configuration for exchanging the GMPLS RSVP-TE signaling messages with the server-layer routers. Routers CE1 and CE2 establish an IP/MPLS adjacency on top of this Layer 2 service.

GMPLS RSVP-TE VLAN LSP Signaling Functionality

Based on [Figure 106 on page 1379](#), the client router establishes the Layer 2 service in the server-layer network as follows:

1. Router CE1 initiates GMPLS RSVP-TE signaling with Router PE1. In this signaling message, Router CE1 indicates the VLAN on the Ethernet link for which it needs the Layer 2 service and the remote CE router, Router CE2, with which the VLAN should be connected.

Router CE1 also indicates the remote PE router, Router PE2, to which Router CE2 is connected, and the exact Ethernet link connecting Router CE2 to Router PE2 on which the Layer 2 service is required in the signaling message.

2. Router PE1 uses the information from Router CE1 in the signaling message and determines the remote PE router, Router PE2, with which Router CE2 is attached. Router PE1 then establishes a packet MPLS LSP (associated bidirectional) through the server-layer MPLS network for carrying the VLAN traffic and then passes the GMPLS RSVP-TE signaling message to Router PE2 using the LSP hierarchy mechanism.

3. Router PE2 propagates the GMPLS RSVP-TE signaling message to Router CE2 with the VLAN to be used on the PE2-CE2 Ethernet link.
4. Router CE2 responds with an acknowledgment to the GMPLS RSVP-TE signaling message to Router PE2. Router PE2 then propagates it to Router PE1, which in turn propagates it to Router CE1.
5. As part of this message propagation, Routers PE1 and PE2 set up the forwarding plane to enable bidirectional flow of VLAN Layer 2 traffic between Routers CE1 and CE2.

LSP Hierarchy with GMPLS RSVP-TE VLAN LSP

The Layer 2 service in GMPLS RSVP-TE VLAN LSP signaling is brought up using a hierarchy mechanism in which two different RSVP LSPs are created for the Layer 2 service:

- An end-to-end VLAN LSP that has state information at the client and server-layer routers.
- An associated bidirectional packet transport LSP that is present in the server-layer routers (PE and P) of the server-layer network.

The LSP hierarchy avoids sharing information about technology-specific LSP characteristics with the core nodes of the server-layer network. This solution cleanly separates the VLAN LSP state and the transport LSP state, and ensures that the VLAN LSP state is only present on the nodes (PE, CE) where it is needed.

Path Specification for GMPLS RSVP-TE VLAN LSP

The path for the GMPLS RSVP-TE LSP is configured as an Explicit Route Object (ERO) at the initiating client router. As this LSP traverses different network domains (initiating, terminating at client network, and traversing the server-layer network), the LSP setup falls under the category of an interdomain LSP setup. In an interdomain scenario, one network domain generally does not have full visibility into the topology of the other network domain. Hence, the ERO that gets configured at the initiating client router does not have full hop information for the server-layer portion. This feature requires that the ERO configured at the CE router has three hops, with the first hop being a strict hop identifying the CE1-PE1 Ethernet link, the second hop being a loose hop identifying the egress PE router (PE2), and the third hop being a strict hop identifying the CE2-PE2 Ethernet link.

GMPLS RSVP-TE VLAN LSP Configuration

The configuration required to set up a GMPLS VLAN LSP at the client and server routers uses the existing GMPLS configuration model with some extensions. The Junos OS GMPLS configuration model for nonpacket LSPs is targeted toward bringing the physical interfaces up and running through GMPLS RSVP-TE signaling, whereas signaling a GMPLS RSVP-TE VLAN LSP aims at bringing up individual VLANs on top of a physical interface. The **ethernet-vlan** configuration statement under the **[edit protocols link-management te-link]** hierarchy enables this.

The client router has physical interfaces connected to a server network, and the server network provides a point-to-point connection between two client routers over the attached physical interfaces. The physical interface is brought into an operational state by GMPLS RSVP-TE as follows:

1. The client router maintains a routing or signaling adjacency with the server network node to which the physical interface is connected, typically through a control channel different from the physical interface, because the physical interface itself is brought up and running only after the signaling.
2. The client router and the server network node identify the physical interfaces connecting them using the TE-link mechanism.
3. The client router and the server network node use the TE-link identifier (IP address) as the GMPLS RSVP hop and the physical interface identifier as the GMPLS label values in the GMPLS RSVP-TE signaling messages to bring the physical interface into an operational state.

In the existing GMPLS configuration, the server and client network nodes use the **protocols link-management peer *peer-name*** configuration statement to specify the adjacent peer node. Because a client router can have one or more physical interfaces connected to the server network node, these physical interfaces are grouped and identified by an IP address through the **protocols link-management te-link *link-name*** configuration statement. The TE-link is assigned a local IP address, a remote IP address, and a list of physical interfaces. The TE-link is then associated with the **protocols link-management peer *peer-name* te-link *te-link-list*** configuration statement.

The out-of-band control channel that is required for exchanging signaling messages is specified using the **protocols link-management peer *peer-name* control-channel *interface-name*** configuration statement. The existence of the server or client network node is made visible to the RSVP and IGP (OSPF) protocols through the **peer-interface *interface-name*** configuration statement under the **[edit protocols rsvp]** and **[edit protocols ospf]** hierarchy levels.

In the existing GMPLS configuration, the label (upstream label and resv label) that is carried in the signaling message is an integer identifier that identifies the physical interface that is required to be brought up. As the label is used to identify the physical interface, the existing GMPLS configuration allows multiple interfaces to be grouped under a single TE-link. In the existing GMPLS configuration, there is sufficient information in the GMPLS RSVP-TE signaling message, such as TE-link address and label value, to identify the physical interface that is required to be brought up. In contrast, for GMPLS RSVP-TE VLAN LSP configuration, the VLAN ID value is used as the label in the signaling message.

In the GMPLS RSVP-TE VLAN LSP configuration, if multiple interfaces are allowed to be configured under a single TE-link, using VLAN ID as the label value in the signaling message can cause ambiguity as to which physical interface on which the VLAN has to be provisioned. Therefore, the TE-link is configured with the **ethernet-vlan** configuration statement, if the number of physical interfaces that can be configured under the TE-link is restricted to only one.

In the existing GMPLS configuration, the bandwidth for a nonpacket LSP is a discrete quantity that corresponds to the bandwidth of the physical interface that needs to be brought up. So, the GMPLS LSP configuration does not allow any bandwidth to be specified, but allows the bandwidth to be specified only

through the **signal-bandwidth** configuration statement under the **[protocols mpls label-switched-path lsp-name lsp-attributes]** hierarchy level. In the GMPLS VLAN LSP configuration, bandwidth is specified similar to that of a packet LSP. In the GMPLS VLAN LSP configuration, the **bandwidth** option is supported and **signal-bandwidth** is not supported.

Associated Bidirectional Packet LSP

The GMPLS RSVP-TE VLAN LSP is carried on an associated bidirectional transport LSP within the server-layer network, which is a single-sided provisioned LSP. The transport LSP signaling is initiated as a unidirectional LSP from the source router to the destination router in the forward direction, and the destination router in turn initiates the signaling of the unidirectional LSP in the reverse direction back to the source router.

Make-Before-Break for Associated Bidirectional Packet and GMPLS RSVP-TE VLAN LSP

The make-before-break support for an associated bidirectional transport LSP follows a similar model, where the destination router for the forward direction of the bidirectional LSP does not perform any make-before-break operations on the reverse direction of the bidirectional LSP. It is the source router (initiator of the associated bidirectional LSP) that initiates the make-before-break newer instance of the associated bidirectional LSP, and the destination router in turn initiates the make-before-break newer instance in the other direction.

For instance, in [Figure 106 on page 1379](#), the unidirectional transport LSP is initiated from Router PE1 to Router PE2 in the forwarding direction, and in turn Router PE2 initiates the transport LSP to Router PE1 in the reverse direction. When a make-before-break instance occurs, only Router PE1 as the initiating client router can establish a new instance of the associated bidirectional LSP. Router PE2 in turn initiates the make-before-break newer instance in the reverse direction.

The make-before-break support for the associated bidirectional transport LSP is used only in scenarios where the transport LSP gets into a state of being locally protected due to link or node failure on the path of the LSP. The GMPLS RSVP-TE VLAN LSP uses the make-before-break mechanism for adjusting seamless bandwidth changes.

NOTE: Periodic re-optimization is not enabled for the associated bidirectional transport LSPs.

The newer make-before-break instance of the GMPLS VLAN LSP is supported under the following constraints:

- It should originate from the same client router as the older instance and be destined to the same client router as the older instance.
- It should use the same server-client links at both the server-client ends as the older instance.

- It should use the same VLAN label at the server-client links as the older instance.
- The GMPLS VLAN LSP should be configured as **adaptive** when the bandwidth change is initiated from the CLI, or else the current instance of the VLAN LSP is torn down and a new VLAN LSP instance is established.

The make-before-break operation for the GMPLS VLAN LSP on the server-layer edge router is rejected if these constraints are not met.

On the server-layer edge routers, when a make-before-break instance of the GMPLS VLAN LSP is seen, a completely new, separate associated bidirectional transport LSP is created to support this make-before-break instance. The existing associated bidirectional LSP (supporting the older instance) is not triggered to initiate a make-before-break instance at the transport LSP level. An implication of this choice (of initiating a new transport LSP) is that at the server-layer resource/bandwidth sharing does not happen when a make-before-break operation is performed for the GMPLS VLAN LSP.

Supported and Unsupported Features

Junos OS supports the following features with the GMPLS RSVP-TE VLAN LSP:

- Request for specific bandwidth and local protection for the VLAN LSP on the client router to the server-layer router.
- Nonstop active routing (NSR) support for the GMPLS VLAN LSP at the client routers, server-layer edge routers, and associated bidirectional transport LSP at the server-layer edge routers.
- Multichassis support.

Junos OS does **not** support the following GMPLS RSVP-TE VLAN LSP functionality:

- Graceful restart support for associated bidirectional packet LSP and GMPLS VLAN LSP.
- End-to-end path computation for GMPLS VLAN LSP using CSPF algorithm at the client router.
- Non-CSPF routing-based discovery of next-hop routers by the different client, server-layer edge routers.
- Automatic provisioning of the client Layer 3 VLAN interfaces upon the successful setup of the VLAN LSP at the client routers.
- MPLS OAM (LSP-ping, BFD).
- Packet MPLS applications, such as next-hop in static route and in IGP shortcuts.
- Local cross connect mechanism, where a client router connects to a remote client router which is connected to the same server router.
- Junos OS Services Framework.
- IPv6 support.
- Logical systems.
- Aggregated Ethernet/SONET/IRB interfaces at the server-client link.

RELATED DOCUMENTATION

[Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling](#) | 1384

Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling

IN THIS SECTION

- [Requirements](#) | 1384
- [Overview](#) | 1385
- [Configuration](#) | 1391
- [Verification](#) | 1405

This example shows how to configure GMPLS RSVP-TE VLAN LSP signaling on the client routers to enable one client router to connect with a remote client router through a server-layer network using the LSP hierarchy. This enables the client routers to establish, maintain, and provision the Layer 2 services, without depending on the server-layer administration, thereby reducing the burden on the operational expenses of the provider network.

Requirements

This example uses the following hardware and software components:

- Six routers that can be a combination of M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, T Series Core Routers, and PTX Series Packet Transport Routers
- Junos OS Release 14.2 or later running on the client routers and server-layer edge routers

Before you begin:

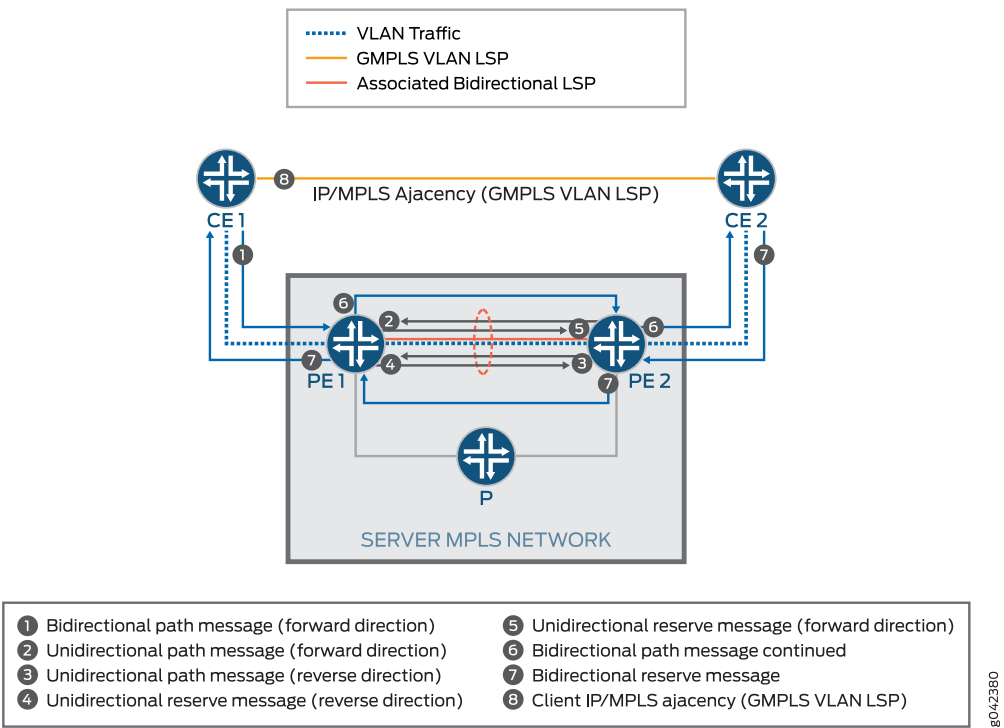
1. Configure the device interfaces.
2. Configure the interface-associated VLANs.
3. Configure the following routing protocols:
 - RSVP
 - MPLS
 - LMP

Overview

Starting with Junos OS Release 14.2, the Layer 2 services between two client routers across an external/third-party server-layer network are set up by the client routers on an on-demand basis through GMPLS RSVP-TE signaling. This feature provides the client routers the flexibility to establish, maintain, and provision the Layer 2 services, without depending on the server-layer administration, thereby reducing the burden on the operational expenses of the provider network. In traditional Layer 2 VPN technology based on LDP and BGP, the provider network handled the provisioning activity for each Layer 2 circuit established between two client routers.

Figure 107 on page 1385 illustrates the setting up and signaling of the GMPLS VLAN LSP between two client routers, CE1 and CE2, across a server-layer network with two server-layer edge routers, PE1 and PE2, and one server-layer core router, P.

Figure 107: Setting Up a GMPLS VLAN LSP



The signaling of GMPLS VLAN LSP is executed as follows:

1. Initiating GMPLS VLAN LSP at CE1

Router CE1 initiates the GMPLS VLAN LSP setup by sending the GMPLS RSVP-TE path message to Router PE1. The signaling between CE1 and PE1 is over an out-of-band control channel, which is a separate control VLAN configured on the Ethernet link connecting the two routers.

The GMPLS RSVP-TE path message initiated by Router CE1 is used to perform the following:

- a. Identify the Ethernet link on which the VLAN is active.
- b. Abstract the Ethernet link as a TE-link and assign an IP address to identify the Ethernet link.
- c. Allocate a VLAN ID from the pool of free VLANs managed by Router CE1 for every Ethernet link connecting Router PE1 to the identified Ethernet link.

This VLAN ID can also be used for the GMPLS VLAN LSP at the CE2-PE2 Ethernet link.

- d. Identify the VLAN for which the Layer 2 service is required to be set up using the allocated VLAN ID as the upstream label object and the upstream direction label value.
- e. Include an ERO object that helps Router PE1 in establishing the VLAN LSP through the server-layer network to the remote client router, CE2. The ERO object in the path message includes three hops:
 - First hop—Strict hop identifying the initiating client-server Ethernet link, PE1-CE1.
 - Second hop—Loose hop identifying the remote server-layer router, PE2.
 - Third hop—Strict hop identifying the remote client-server Ethernet link, PE2-CE2.
- f. Include the bandwidth required for the GMPLS VLAN LSP.
- g. Include any local-protection required within the server-layer network for the VLAN LSP.

2. Initiating Associated Bidirectional Transport LSP at PE1

After Router PE1 receives the path message from Router CE1, the message is validated to check the availability of the Ethernet link and VLAN ID. In the server-layer network, the Layer 2 services between the server-layer routers, PE1 and PE2, are provided at the data plane in a manner similar to Layer 2 circuits. Router PE1 brings up a transport LSP to Router PE2 and then extends the GMPLS VLAN LSP as a hierarchical LSP running on top of the PE1-PE2 transport LSP. The PE1-PE2 transport LSP is a packet LSP and is bidirectional in nature. This is because the GMPLS VLAN LSP is bidirectional and each server-layer router needs to be able to do the following:

- Receive traffic from the server-client Ethernet link (for example, the PE1-CE1 link) and send it to the remote server-layer router, PE2.
- Receive traffic from remote Router PE2 and send it on the PE1-CE1 Ethernet link.

For each GMPLS VLAN LSP, a packet transport LSP is set up within the server-layer network. The transport LSP is exclusively used to carry traffic of the GMPLS VLAN LSP for which it was created. The

transport LSP is dynamically created at the time of receiving the GMPLS VLAN LSP; thus, no configuration is required to trigger its creation. The transport LSP established for the VLAN LSP inherits the bandwidth and the local-protection attributes from the VLAN LSP.

Router PE1 signals the PE1-PE2 transport LSP to Router PE2. Router PE1 determines the destination for the transport LSP from the loose hop specified in the ERO object of the GMPLS RSVP-TE path message from Router CE1 and then signals the VLAN LSP. However, if the PE1-PE2 transport LSP fails to establish, Router PE1 sends back a path error message to Router CE1, and the GMPLS VLAN LSP is not established as well.

3. Setting Up the Associated Bidirectional Transport LSP Between the Server-Layer Routers

The associated bidirectional LSP between routers PE1 and PE2 consists of two unidirectional packet LSPs:

- PE1-to-PE2
- PE2-to-PE1

Router PE1 initiates signaling of a unidirectional packet LSP to Router PE2. This unidirectional packet LSP constitutes the forward direction (PE1-to-PE2) of the associated bidirectional LSP, and the path message carries the Extended Association Object indicating this is a single-sided provisioning model. On receiving the path message for the LSP, Router PE2 responds with a Resv message and triggers the signaling of a unidirectional packet LSP to Router PE1 with the same path as (PE1-to-PE2) in the reverse direction. This unidirectional packet LSP uses the PE2-to-PE1 direction of the associated bidirectional LSP, and this path message carries the same Extended Association Object seen in the PE1-to-PE2 path message.

When Router PE1 receives the Resv message for the PE1-to-PE2 unidirectional LSP and the path message for the PE2-to-PE1 unidirectional LSP, PE1 binds the PE1-to-PE2 and PE2-to-PE1 unidirectional LSPs by matching the Extended Association Objects carried in the respective path messages. For the path message for the PE2-to-PE1 unidirectional LSP, Router PE1 responds with the Resv Message. On receiving the Resv message for the PE1-to-PE2 LSP and the path message for the PE2-to-PE1 LSP, Router PE1 has established the associated bidirectional packet transport LSP.

4. Setting Up the GMPLS VLAN LSP at Router PE1

After successfully establishing the transport LSP, Router PE1 triggers the signaling of the GMPLS VLAN LSP. Router PE1 sends the GMPLS RSVP-TE path message corresponding to the VLAN LSP directly to Router PE2, which is bidirectional in nature and includes the upstream label object.

Router PE2 is not aware of the association between the transport LSP and the VLAN LSP. This association is indicated to Router PE2 by Router PE1.

5. Setting Up the GMPLS VLAN LSP at Router PE2

On receiving the VLAN LSP path message from Router PE1, Router PE2 verifies the availability of the transport LSP. If the transport LSP is not available or the LSP setup is in progress, the VLAN LSP processing is put on hold. When the transport LSP is available, Router PE2 processes the VLAN LSP path message. The ERO object in this path message indicates that the next hop is a strict hop identifying the PE2-to-CE2 Ethernet link. The ERO object can indicate the VLAN ID to be used on the PE2-to-CE2 Ethernet link by Router PE2.

Router PE2 appropriately allocates the VLAN ID to be sent as the upstream label in the VLAN LSP path message to Router CE2, and sends it through an out-of-band control channel.

6. Processing the GMPLS VLAN LSP at Router CE2

On receiving the GMPLS RSVP-TE LSP from Router PE2, Router CE2 validates the availability of VLAN ID for allocation on the PE2-to-CE2 link. Router CE2 then allocates the VLAN ID for this VLAN LSP and sends back a Resv message to Router PE2 with the VLAN ID as the label object in the Resv message.

7. Processing the GMPLS VLAN LSP at Router PE2

On receiving the Resv message from Router CE2, Router PE2 validates that the label object in the Resv message has the same VLAN ID as in the path message. Router PE2 then allocates a 20-bit MPLS label, which is included in the Resv message sent to Router PE1.

Router PE2 then programs the forwarding plane with the entries to provide the Layer 2 service functionality.

NOTE: For all the VLAN IDs that can be allocated as labels on the PE1-to-CE1 and PE2-CE2 Ethernet links, you must manually configure logical interfaces for circuit cross-connect (CCC) purposes on the server-layer edge routers and not for other families, such as IPv4, IPv6, or MPLS.

8. Processing the GMPLS VLAN LSP at Router PE1

On receiving the Resv message for the VLAN LSP from Router PE2, Router PE1 sends a Resv message to Router CE1 with the same VLAN ID it received as the upstream label from Router CE1. Router PE1 programs the forwarding plane with the entries to provide the Layer 2 service functionality as Router PE2.

9. Processing the GMPLS VLAN LSP at Router CE1

On receiving the Resv message from Router PE1, Router CE1 validates that the VLAN ID received in the Resv message matches the VLAN ID in the upstream label in the path message it sent. This completes the setup of the GMPLS VLAN LSP from Router CE1 to Router CE2.

NOTE:

- The GMPLS VLAN LSP setup does not result in the addition of any forwarding plane entries at the client routers, CE1 and CE2. Only the server-layer routers, PE1 and PE2, add the forwarding plane entries for the GMPLS VLAN LSP.
- There is no routing information exchange between the client and the server-layer routers. The client and server-layer routers do not exchange their network topology information with each other.

10. Accounting for Bandwidth of the GMPLS VLAN LSP

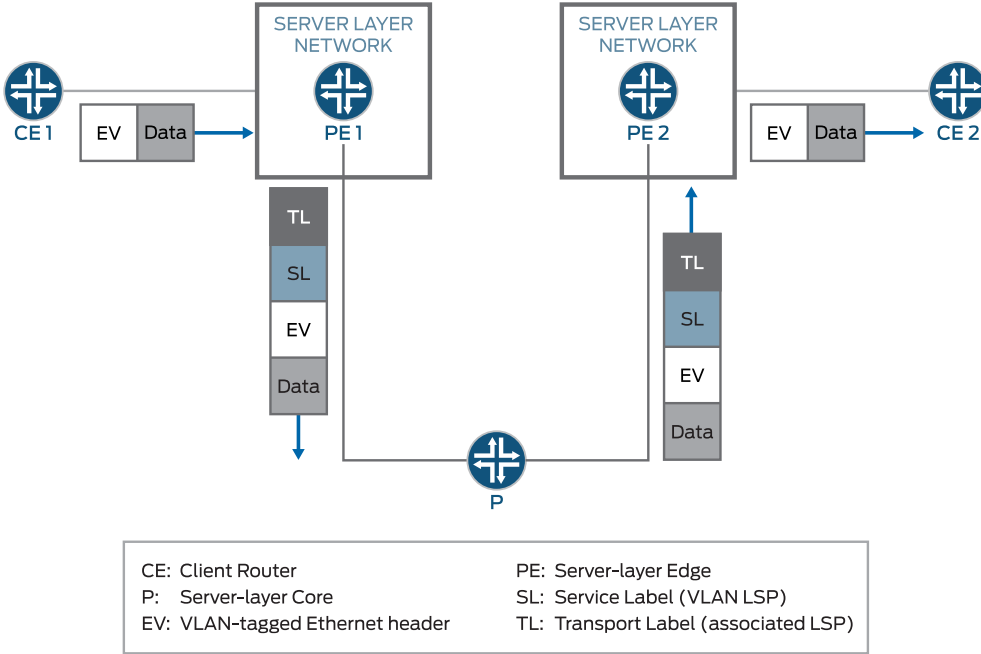
On successfully setting up the GMPLS VLAN LSP, both the client and server-layer routers reduce the amount of available bandwidth on the server-client Ethernet links by the bandwidth amount allocated for the GMPLS VLAN LSP. This bandwidth accounting information is used for admission control purposes when additional GMPLS VLAN LSPs are brought up on the server-client Ethernet links.

11. Using GMPLS VLAN LSP by the Client Routers

After successfully setting up the GMPLS VLAN LSP, the client routers – CE1 and CE2 – need to be manually configured with the VLAN logical interface on top of the server-client Ethernet links with the signaled VLAN ID. This logical interface needs to be configured with the IP address and needs to be included in the IGP protocol. As a result of this configuration, Routers CE1 and CE2 establish IGP adjacency and exchange data traffic over the Layer 2 service established through the GMPLS signaling.

[Figure 108 on page 1390](#) illustrates the data traffic flow of the GMPLS VLAN LSP from Router CE1 to Router CE2 after the LSP setup is complete and the necessary CE1-to-CE2 IGP/MPLS adjacency has been established. The server-layer transport LSP originates from Router PE1, traverses a single server-layer core router, Router P, and reaches Router PE2. The server-layer transport LSP is shown as a penultimate-hop pop LSP, where Router P pops off the transport LSP label, and only the service label is present on the P-to-PE2 link.


Figure 108: Data Traffic Flow of GMPLS VLAN LSP



Topology

In [Figure 109 on page 1390](#), GMPLS RSVP-TE VLAN LSP signaling is used to establish the Layer 2 services between the client routers, Router CE1 and Router CE2. The server routers, Router PE1 and Router PE2, have a GRE tunnel established with each of the directly connected client routers. Routers P1 and P2 are also server routers in the server-layer network.

Figure 109: Configuring GMPLS RSVP-TE VLAN LSP Signaling



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

CE1

```

set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 1 family inet address 1.1.1.1/30
set interfaces ge-0/0/0 unit 1 family mpls
set interfaces ge-0/0/0 unit 10 vlan-id 10
set interfaces ge-0/0/0 unit 10 family inet address 10.10.10.1/24
set interfaces ge-0/0/0 unit 10 family mpls
set interfaces gre unit 0 tunnel source 1.1.1.1
set interfaces gre unit 0 tunnel destination 1.1.1.2
set interfaces gre unit 0 family inet address 10.35.100.25/30
set interfaces lo0 unit 0 family inet address 10.255.10.1/32
set routing-options router-id 10.255.10.1
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp peer-interface PE1
set protocols mpls no-cspf
set protocols mpls label-switched-path CE1-to-CE2 from 10.255.10.1
set protocols mpls label-switched-path CE1-to-CE2 to 10.255.10.6
set protocols mpls label-switched-path CE1-to-CE2 lsp-attributes switching-type ethernet-vlan
set protocols mpls label-switched-path CE1-to-CE2 lsp-attributes upstream-label vlan-id 10
set protocols mpls label-switched-path CE1-to-CE2 bandwidth 100m
set protocols mpls label-switched-path CE1-to-CE2 primary path1
set protocols mpls path path1 10.35.1.2 strict
set protocols mpls path path1 10.255.10.5 loose
set protocols mpls path path1 10.36.1.1 strict
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols link-management te-link link10 local-address 10.35.1.1
set protocols link-management te-link link10 remote-address 10.35.1.2
set protocols link-management te-link link10 ethernet-vlan
set protocols link-management te-link link10 interface ge-0/0/0
set protocols link-management peer PE1 address 10.255.10.2
set protocols link-management peer PE1 control-channel gre.0
set protocols link-management peer PE1 te-link link10

```


PE1

```

set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 encapsulation flexible-ethernet-services
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 1 family inet address 1.1.1.2/30
set interfaces ge-0/0/0 unit 1 family mpls
set interfaces ge-0/0/0 unit 10 encapsulation vlan-ccc
set interfaces ge-0/0/0 unit 10 vlan-id 10
set interfaces ge-0/0/1 unit 0 family inet address 70.70.70.1/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces gre unit 0 tunnel source 1.1.1.2
set interfaces gre unit 0 tunnel destination 1.1.1.1
set interfaces gre unit 0 family inet address 10.35.100.26/30
set interfaces lo0 unit 0 family inet address 10.255.10.2/32
set routing-options router-id 10.255.10.2
set protocols rsvp associated-bidirectional-lsp single-sided-provisioning
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp peer-interface CE1 dynamic-bidirectional-transport
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols link-management te-link link1 local-address 10.35.1.2
set protocols link-management te-link link1 remote-address 10.35.1.1
set protocols link-management te-link link1 ethernet-vlan vlan-id-range 1-1000
set protocols link-management te-link link1 interface ge-0/0/0
set protocols link-management peer CE1 address 10.255.10.1
set protocols link-management peer CE1 control-channel gre.0
set protocols link-management peer CE1 te-link link1

```

P1

```

set interfaces ge-0/0/0 unit 0 family inet address 90.90.90.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 70.70.70.2/24

```



```

set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 80.80.80.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.10.3/32
set routing-options router-id 10.255.10.3
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

P2

```

set interfaces ge-0/0/0 unit 0 family inet address 90.90.90.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 30.30.30.1/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.2/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.10.4/32
set routing-options router-id 10.255.10.4
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

PE2

```

set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 encapsulation flexible-ethernet-services
set interfaces ge-0/0/0 unit 0 vlan-id 1
set interfaces ge-0/0/0 unit 0 family inet address 2.2.2.2/30

```



```

set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/0 unit 10 encapsulation vlan-ccc
set interfaces ge-0/0/0 unit 10 vlan-id 10
set interfaces ge-0/0/1 unit 0 family inet address 30.30.30.2/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 80.80.80.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces gre unit 0 tunnel source 2.2.2.2
set interfaces gre unit 0 tunnel destination 2.2.2.1
set interfaces gre unit 0 family inet address 10.35.101.26/30
set interfaces lo0 unit 0 family inet address 10.255.10.5/32
set routing-options router-id 10.255.10.5
set protocols rsvp associated-bidirectional-lsp single-sided-provisioning
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp peer-interface CE2 dynamic-bidirectional-transport
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols link-management te-link link1 local-address 10.36.1.2
set protocols link-management te-link link1 remote-address 10.36.1.1
set protocols link-management te-link link1 ethernet-vlan vlan-id-range 1-1000
set protocols link-management te-link link1 interface ge-0/0/0
set protocols link-management peer CE2 address 10.255.10.6
set protocols link-management peer CE2 control-channel gre.0
set protocols link-management peer CE2 te-link link1

```

CE2

```

set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 1 family inet address 2.2.2.1/24
set interfaces ge-0/0/0 unit 1 family mpls
set interfaces ge-0/0/0 unit 10 vlan-id 10
set interfaces ge-0/0/0 unit 10 family inet address 10.10.10.2/24
set interfaces ge-0/0/0 unit 10 family mpls
set interfaces gre unit 0 tunnel source 2.2.2.1
set interfaces gre unit 0 tunnel destination 2.2.2.2

```



```

set interfaces gre unit 0 family inet address 10.35.101.25/30
set interfaces lo0 unit 0 family inet address 10.255.10.6/32
set routing-options router-id 10.255.10.6
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp peer-interface PE2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols link-management te-link link10 local-address 10.36.1.1
set protocols link-management te-link link10 remote-address 10.36.1.2
set protocols link-management te-link link10 ethernet-vlan vlan-id-range 1-1000
set protocols link-management te-link link10 interface ge-0/0/0
set protocols link-management peer PE2 address 10.255.10.5
set protocols link-management peer PE2 control-channel gre.0
set protocols link-management peer PE2 te-link link10

```

Configuring the Client Router

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE1:

NOTE: Repeat this procedure for Router CE2 in the server-layer network, after modifying the appropriate interface names, addresses, and any other parameters for the router.

1. Configure the interface connecting Router CE1 to Router PE1.

```

[edit interfaces]
user@CE1# set ge-0/0/0 vlan-tagging

```

2. Configure the control VLAN for the ge-0/0/0 interface.

```

[edit interfaces]
user@CE1# set ge-0/0/0 unit 1 vlan-id 1
user@CE1# set ge-0/0/0 unit 1 family inet address 1.1.1.1/30
user@CE1# set ge-0/0/0 unit 1 family mpls

```


3. Configure the LSP VLAN on the ge-0/0/0 interface.

```
[edit interfaces]
user@CE1# set ge-0/0/0 unit 10 vlan-id 10
user@CE1# set ge-0/0/0 unit 10 family inet address 10.10.10.1/24
user@CE1# set ge-0/0/0 unit 10 family mpls
```

4. Configure the GRE tunnel as the controlling interface for Router CE1.

```
[edit interfaces]
user@CE1# set gre unit 0 tunnel source 1.1.1.1
user@CE1# set gre unit 0 tunnel destination 1.1.1.2
user@CE1# set gre unit 0 family inet address 10.35.100.25/30
```

5. Configure the loopback interface of Router CE1.

```
[edit interfaces]
user@CE1# set lo0 unit 0 family inet address 10.255.10.1/32
```

6. Configure the loopback address of Router CE1 as its router ID.

```
[edit routing-options]
user@CE1# set router-id 10.255.10.1
```

7. Enable RSVP on all the interfaces of Router CE1, excluding the management interface.

```
[edit protocols]
user@CE1# set rsvp interface all
user@CE1# set rsvp interface fxp0.0 disable
```

8. Configure the RSVP peer interface for Router CE1.

```
[edit protocols]
user@CE1# set rsvp peer-interface PE1
```

9. Disable automatic path computation for label-switched paths (LSPs).

```
[edit protocols]
```



```
user@CE1# set mpls no-cspf
```

10. Configure the LSP to connect Router CE1 to Router CE2.

```
[edit protocols]
user@CE1# set mpls label-switched-path CE1-to-CE2 from 10.255.10.1
user@CE1# set mpls label-switched-path CE1-to-CE2 to 10.255.10.6
```

11. Configure the CE1-to-CE2 LSP attributes.

```
[edit protocols]
user@CE1# set mpls label-switched-path CE1-to-CE2 lsp-attributes switching-type ethernet-vlan
user@CE1# set mpls label-switched-path CE1-to-CE2 lsp-attributes upstream-label vlan-id 10
user@CE1# set mpls label-switched-path CE1-to-CE2 bandwidth 100m
```

12. Configure the CE1-to-CE2 LSP path and path parameters.

```
[edit protocols]
user@CE1# set mpls label-switched-path CE1-to-CE2 primary path1
user@CE1# set mpls path path1 10.35.1.2 strict
user@CE1# set mpls path path1 10.255.10.5 loose
user@CE1# set mpls path path1 10.36.1.1 strict
```

13. Enable MPLS on all the interfaces of Router CE1, excluding the management interface.

```
[edit protocols]
user@CE1# set mpls interface all
user@CE1# set mpls interface fxp0.0 disable
```

14. Configure a traffic engineering link, and assign addresses for the local and remote end of the link.

```
[edit protocols]
user@CE1# set link-management te-link link10 local-address 10.35.1.1
user@CE1# set link-management te-link link10 remote-address 10.35.1.2
```

15. Enable setting up of Layer 2 VLAN LSP on the link10 traffic engineering link.

```
[edit protocols]
```



```
user@CE1# set link-management te-link link10 ethernet-vlan
```

16. Configure the Router CE1 interface as the member interface of the link10 traffic engineering link.

```
[edit protocols]
user@CE1# set link-management te-link link10 interface ge-0/0/0
```

17. Configure Router PE1 as the Link Management Protocol (LMP) peer for Router CE1, and configure the peer attributes.

```
[edit protocols]
user@CE1# set link-management peer PE1 address 10.255.10.2
user@CE1# set link-management peer PE1 control-channel gre.0
user@CE1# set link-management peer PE1 te-link link10
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE1# show interfaces
ge-0/0/0 {
  vlan-tagging;
  unit 1 {
    vlan-id 1;
    family inet {
      address 1.1.1.1/30;
    }
    family mpls;
  }
  unit 10 {
    vlan-id 10;
    family inet {
      address 10.10.10.1/24;
    }
    family mpls;
  }
}
gre {
  unit 0 {
```



```

    tunnel {
        source 1.1.1.1;
        destination 1.1.1.2;
    }
    family inet {
        address 10.35.100.25/30;
    }
}
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.10.1/32;
        }
    }
}
}

```

```

user@CE1# show routing-options
router-id 10.255.10.1;

```

```

user@CE1# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
    peer-interface PE1;
}
mpls {
    no-cspf;
    label-switched-path CE1-to-CE2 {
        from 10.255.10.1;
        to 10.255.10.6;
        lsp-attributes {
            switching-type ethernet-vlan;
            upstream-label {
                vlan-id 10;
            }
        }
        bandwidth 100m;
        primary path1;
    }
    path path1 {

```



```

    10.35.1.2 strict;
    10.255.10.5 loose;
    10.36.1.1 strict;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
}
link-management {
  te-link link10 {
    local-address 10.35.1.1;
    remote-address 10.35.1.2;
    ethernet-vlan;
    interface ge-0/0/0;
  }
  peer PE1 {
    address 10.255.10.2;
    control-channel gre.0;
    te-link link10;
  }
}
}

```

Configuring the Server Router

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

NOTE: Repeat this procedure for Router PE2 in the server-layer network, after modifying the appropriate interface names, addresses, and any other parameters for the router.

1. Configure the interface connecting Router PE1 to Router CE1.

```

[edit interfaces]
user@PE1# set ge-0/0/0 vlan-tagging
user@PE1# set ge-0/0/0 encapsulation flexible-ethernet-services

```

2. Configure the control VLAN for the ge-0/0/0 interface.


```
[edit interfaces]
user@PE1# set ge-0/0/0 unit 1 vlan-id 1
user@PE1# set ge-0/0/0 unit 1 family inet address 1.1.1.2/30
user@PE1# set ge-0/0/0 unit 1 family mpls
```

3. Configure the LSP VLAN on the ge-0/0/0 interface.

```
[edit interfaces]
user@PE1# set ge-0/0/0 unit 10 encapsulation vlan-ccc
user@PE1# set ge-0/0/0 unit 10 vlan-id 10
```

4. Configure the interface connecting Router PE1 to the core routers (Router P1 and Router P2).

```
[edit interfaces]
user@PE1# set ge-0/0/1 unit 0 family inet address 70.70.70.1/30
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set ge-0/0/2 unit 0 family inet address 20.20.20.1/30
user@PE1# set ge-0/0/2 unit 0 family mpls
```

5. Configure the GRE tunnel as the controlling interface for Router PE1.

```
[edit interfaces]
user@PE1# set gre unit 0 tunnel source 1.1.1.2
user@PE1# set gre unit 0 tunnel destination 1.1.1.1
user@PE1# set gre unit 0 family inet address 10.35.100.26/30
```

6. Configure the loopback interface of Router PE1.

```
[edit interfaces]
user@PE1# set lo0 unit 0 family inet address 10.255.10.2/32
```

7. Configure the loopback address of Router PE1 as its router ID.

```
[edit routing-options]
user@PE1# set router-id 10.255.10.2
```

8. Configure an associated bidirectional LSP, and enable unidirectional reverse LSP setup for single-sided provisioned forward LSP.


```
[edit protocols]
user@PE1# set rsdp associated-bidirectional-lsp single-sided-provisioning
```

9. Enable RSVP on all the interfaces of Router PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set rsdp interface all
user@PE1# set rsdp interface fxp0.0 disable
```

10. Configure the RSVP peer interface for Router PE1, and enable dynamic setup of bidirectional packet LSP for transporting nonpacket GMPLS LSP.

```
[edit protocols]
user@PE1# set rsdp peer-interface CE1 dynamic-bidirectional-transport
```

11. Enable MPLS on all the interfaces of Router PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
```

12. Configure OSPF with traffic engineering capabilities.

```
[edit protocols]
user@PE1# set ospf traffic-engineering
```

13. Enable OSPF area 0 on all the interfaces of Router PE1, excluding the management interface.

```
[edit protocols]
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
```

14. Configure a traffic engineering link, and assign addresses for the local and remote end of the link.

```
[edit protocols]
user@PE1# set link-management te-link link1 local-address 10.35.1.2
user@PE1# set link-management te-link link1 remote-address 10.35.1.1
```


15. Enable setting up of a Layer 2 VLAN LSP for a specific range of VLANs on the link1 traffic engineering link.

```
[edit protocols]
user@PE1# set link-management te-link link1 ethernet-vlan vlan-id-range 1-1000
```

16. Configure the Router PE1 interface as the member interface of the link1 traffic engineering link.

```
[edit protocols]
user@CE1# set link-management te-link link1 interface ge-0/0/0
```

17. Configure Router CE1 as the LMP peer for Router PE1, and configure the peer attributes.

```
[edit protocols]
user@CE1# set link-management peer CE1 address 10.255.10.1
user@CE1# set link-management peer CE1 control-channel gre.0
user@CE1# set link-management peer CE1 te-link link1
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
ge-0/0/0 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 1 {
    vlan-id 1;
    family inet {
      address 1.1.1.2/30;
    }
    family mpls;
  }
  unit 10 {
    encapsulation vlan-ccc;
    vlan-id 10;
  }
}
ge-0/0/1 {
  unit 0 {
```



```

        family inet {
            address 70.70.70.1/30;
        }
        family mpls;
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 20.20.20.1/30;
        }
        family mpls;
    }
}
gre {
    unit 0 {
        tunnel {
            source 1.1.1.2;
            destination 1.1.1.1;
        }
        family inet {
            address 10.35.100.26/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.10.2/32;
        }
    }
}

```

```

user@PE1# show routing-options
router-id 10.255.10.2;

```

```

user@PE1# show protocols
rsvp {
    associated-bidirectional-lsp single-sided-provisioning;
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```



```

peer-interface CE1 {
    dynamic-bidirectional-transport;
}
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
link-management {
    te-link link1 {
        local-address 10.35.1.2;
        remote-address 10.35.1.1;
        ethernet-vlan {
            vlan-id-range 1-1000;
        }
        interface ge-0/0/0;
    }
    peer CE1 {
        address 10.255.10.1;
        control-channel gre.0;
        te-link link1;
    }
}
}

```

Verification

IN THIS SECTION

- [Verifying the Traffic Engineering Link Status on the Client Routers | 1406](#)
- [Verifying the RSVP Session Status on the Client Routers | 1407](#)
- [Verifying the LSP Status on the Server Router | 1408](#)

- Verifying the CCC Entries in the MPLS Routing Table of the Server Routers | 1409
- Verifying End-to-End Connectivity | 1410

Confirm that the configuration is working properly.

Verifying the Traffic Engineering Link Status on the Client Routers

Purpose

Verify the status of the traffic engineering link configured between Router CE1 and Router CE2.

Action

From operational mode, run the **show link-management** and the **show link-management te-link detail** commands.

```
user@CE1> show link-management
```

```
Peer name: PE1, System identifier: 50740
State: Up, Control address: 10.255.10.2
Hello interval: 150, Hello dead interval: 500
  Control-channel      State
  gre.0                Active
TE links:
  link10

TE link name: link10, State: Up
  Local identifier: 65075, Remote identifier: 0, Local address: 10.35.1.1, Remote
address: 10.35.1.2, Encoding: Ethernet, Switching: EVPL, Minimum bandwidth: 0bps,

  Maximum bandwidth: 1000Mbps, Total bandwidth: 1000Mbps, Available bandwidth:
900Mbps
  Name                  State Local ID Remote ID      Bandwidth Used
  LSP-name
  ge-0/0/0              Up      54183      0      1000Mbps Yes
  CE1-to-CE2
```

```
user@CE1> show link-management te-link detail
```

```
TE link name: link10, State: Up
  Local identifier: 65075, Remote identifier: 0, Local address: 10.35.1.1, Remote
```



```

address: 10.35.1.2, Encoding: Ethernet, Switching: EVPL, Minimum bandwidth: 0bps,

Maximum bandwidth: 1000Mbps, Total bandwidth: 1000Mbps, Available bandwidth:
900Mbps
Resource: ge-0/0/0, Type: IFD, System identifier: 137, State: Up, Local
identifier: 54183, Remote identifier: 0
Total bandwidth: 1000Mbps, Unallocated bandwidth: 900Mbps
Traffic parameters: Encoding: Ethernet, Switching: EVPL, Granularity: Unknown

Maximum allocations: 4094, Number of allocations: 1, Unique allocations: 1,
In use: Yes
LSP name: CE1-to-CE2, Local label: 10, Remote label: 10, Allocated bandwidth:
100Mbps

```

user@CE2> show link-management

```

Peer name: PE2, System identifier: 50743
State: Up, Control address: 10.255.10.5
Hello interval: 150, Hello dead interval: 500
Control-channel          State
gre.0                    Active
TE links:
link10

TE link name: link10, State: Up
Local identifier: 65075, Remote identifier: 0, Local address: 10.36.1.1, Remote
address: 10.36.1.2, Encoding: Ethernet, Switching: EVPL, Minimum bandwidth: 0bps,

Maximum bandwidth: 1000Mbps, Total bandwidth: 1000Mbps, Available bandwidth:
900Mbps

```

Name	State	Local ID	Remote ID	Bandwidth	Used
LSP-name					
ge-0/0/0	Up	54183	0	1000Mbps	Yes
CE1-to-CE2					

Meaning

The Link Management Protocol (LMP) peering has been established between the client routers, and the traffic engineering link is up on both Routers CE1 and CE2.

Verifying the RSVP Session Status on the Client Routers

Purpose

Verify the status of the RSVP sessions between Router CE1 and Router CE2.

Action

From operational mode, run the **show rsvp session** command.

```
user@CE1> show rsvp session
```

Ingress RSVP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.6	10.255.10.1	Up	0	1 FF	-	10	CE1-to-CE2 Bidir

Total 1 displayed, Up 1, Down 0

Egress RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

```
user@CE2> show rsvp session
```

Ingress RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

Egress RSVP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.6	10.255.10.1	Up	0	1 FF	10	-	CE1-to-CE2 Bidir

Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

Meaning

The RSVP sessions are established between the ingress router, Router CE1, and the egress router, Router CE2.

Verifying the LSP Status on the Server Router

Purpose

Verify the status of the MPLS LSP on Router PE1.

Action

From operational mode, run the **show mpls lsp** command.


```
user@PE1> show mpls lsp
```

```
Ingress LSP: 1 sessions
To          From          State Rt P    ActivePath    LSPname
10.255.10.5  10.255.10.2    Up    0  *
vlan:0:10:8176:10.255.10.2->10.255.10.5 Assoc-Bidir
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.255.10.2  10.255.10.5    Up    0  1 FF      3      -
vlan:0:10:8176:10.255.10.2->10.255.10.5:rev
Total 1 displayed, Up 1, Down 0

Transit LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.255.10.6  10.255.10.1    Up    0  1 FF      10     299808 CE1-to-CE2 Bidir
Total 1 displayed, Up 1, Down 0
```

Meaning

The CE1-to-CE2 LSP is established, and the output displays the LSP attributes.

Verifying the CCC Entries in the MPLS Routing Table of the Server Routers

Purpose

Verify the circuit cross-connect (CCC) interface entries in the MPLS routing table.

Action

From operational mode, run the **show route table mpls.0** and the **show route forwarding-table ccc ccc-interface** commands.

```
user@PE1> show route table mpls.0
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 1d 22:14:51, metric 1
           Receive
1          *[MPLS/0] 1d 22:14:51, metric 1
           Receive
2          *[MPLS/0] 1d 22:14:51, metric 1
           Receive
13         *[MPLS/0] 1d 22:14:51, metric 1
```



```

                Receive
299824          *[RSVP/7/1] 17:32:07, metric 1
                > via ge-0/0/0.10, Pop
ge-0/0/0.10     *[RSVP/7/1] 17:32:07, metric 1
                > to 20.20.20.2 via ge-0/0/2.0, label-switched-path CE1-to-CE2

```

```
user@PE1> show route forwarding-table ccc ge-0/0/0.10
```

```

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index   NhRef Netif
ge-0/0/0.10      (CCC) user    0 20.20.20.2      Push 299808, Push 299872(top)
581              2 ge-0/0/2.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index   NhRef Netif
default          perm    0                      dscd    534     1

```

Meaning

The output displays the CCC interface that is the client-router-facing interface and the next-hop details for that interface.

Verifying End-to-End Connectivity

Purpose

Verify the connectivity between Router CE1 and the remote client router, Router CE2.

Action

From operational mode, run the **ping** command.

```
user@CE1> ping 10.10.10.2
```

```

PING 10.10.10.2 (10.10.10.2): 56 data bytes
64 bytes from 10.10.10.2: icmp_seq=0 ttl=64 time=15.113 ms
64 bytes from 10.10.10.2: icmp_seq=1 ttl=64 time=13.353 ms
64 bytes from 10.10.10.2: icmp_seq=2 ttl=64 time=13.769 ms
64 bytes from 10.10.10.2: icmp_seq=3 ttl=64 time=10.341 ms
64 bytes from 10.10.10.2: icmp_seq=4 ttl=64 time=12.597 ms
^C
--- 10.10.10.2 ping statistics ---

```



```
5 packets transmitted, 5 packets received, 0% packet loss  
round-trip min/avg/max/stddev = 10.341/13.035/15.113/1.575 ms
```

Meaning

The ping from Router CE1 to Router CE2 is successful.

RELATED DOCUMENTATION

[GMPLS RSVP-TE VLAN LSP Signaling Overview](#) | 1377

7

PART

MPLS BGP VPNs

Configuring MPLS VPNs | **1415**

Configuring TCC and Layer 2.5 Switching | **1515**

Configuring CLNS VPNs | **1521**

Configuring VPLS | **1545**

Configuring Circuit Cross-Connect (CCC) and Translational Cross-Connect (TCC) | **1643**

Configuring MPLS VPNs

IN THIS CHAPTER

- [MPLS VPN Overview | 1416](#)
- [Understanding IPv6 Layer 3 VPNs | 1420](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on Switches | 1420](#)
- [Configuring a BGP Session for MPLS VPNs \(CLI Procedure\) | 1421](#)
- [Configuring an IGP and the RSVP Signaling Protocol \(CLI Procedure\) | 1423](#)
- [Configuring Routing Options for MPLS VPNs \(CLI Procedure\) | 1423](#)
- [Configuring a Routing Instance for MPLS VPNs \(CLI Procedure\) | 1424](#)
- [Chained Composite Next Hops for Transit Devices for VPNs | 1426](#)
- [Understanding MPLS Layer 2 VPNs | 1427](#)
- [Understanding Ethernet-over-MPLS \(L2 Circuit\) | 1428](#)
- [MPLS Layer 2 VPN Configuration Overview | 1429](#)
- [Configuring a Routing Policy for MPLS Layer 2 VPNs \(CLI Procedure\) | 1430](#)
- [Configuring Interfaces for Layer 2 VPNs \(CLI Procedure\) | 1432](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) | 1433](#)
- [Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)
- [Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches | 1456](#)
- [Verifying an MPLS Layer 2 VPN Configuration | 1459](#)
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)
- [Understanding MPLS Layer 2 Circuits | 1463](#)
- [MPLS Layer 2 Circuit Configuration Overview | 1464](#)
- [Configuring an MPLS Layer 2 Circuit \(CLI Procedure\) | 1465](#)
- [Verifying an MPLS Layer 2 Circuit Configuration | 1466](#)
- [Configuring an IGP and the LDP Signaling Protocol \(CLI Procedure\) | 1466](#)
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1468](#)
- [Understanding MPLS Layer 3 VPNs | 1471](#)
- [MPLS Layer 3 VPN Configuration Overview | 1471](#)
- [Configuring a Routing Policy for MPLS Layer 3 VPNs \(CLI Procedure\) | 1473](#)
- [Verifying an MPLS Layer 3 VPN Configuration | 1474](#)

- [Example: Configuring MPLS-Based Layer 3 VPNs | 1474](#)
- [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks | 1486](#)
- [Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches | 1498](#)
- [Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) | 1512](#)

MPLS VPN Overview

IN THIS SECTION

- [MPLS VPN Topology | 1416](#)
- [MPLS VPN Routing | 1418](#)
- [VRF Instances | 1418](#)
- [Route Distinguishers | 1418](#)
- [MPLS Exceptions on SRX Series Devices | 1419](#)

Virtual private networks (VPNs) are private networks that use a public network to connect two or more remote sites. Instead of dedicated connections between networks, VPNs use virtual connections routed (tunneled) through public networks that are typically service provider networks. VPNs are a cost-effective alternative to expensive dedicated lines. The type of VPN is determined by the connections it uses and whether the customer network or the provider network performs the virtual tunneling.


You can configure a router running Junos OS to participate in several types of VPNs. This topic discusses MPLS VPNs.

This topic contains the following sections:

MPLS VPN Topology

There are many ways to set up an MPLS VPN and direct traffic through it. [Figure 110 on page 1417](#) shows a typical MPLS VPN topology.

Figure 110: Typical VPN Topology



There are three primary types of MPLS VPNs: Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs. All types of MPLS VPNs share certain components:

- The provider edge (PE) routers in the provider's network connect to the customer edge (CE) routers located at customer sites. PE routers support VPN and MPLS label functionality. Within a single VPN, pairs of PE routers are connected through a virtual tunnel, typically a label-switched path (LSP).
- Provider routers within the core of the provider's network are not connected to any routers at a customer site but are part of the tunnel between pairs of PE routers. Provider routers support LSP functionality as part of the tunnel support, but do not support VPN functionality.
- CE routers are the routers or switches located at the customer site that connect to the provider's network. CE routers are typically IP routers, but they can also be Asynchronous Transfer Mode (ATM), Frame Relay, or Ethernet switches.

All VPN functions are performed by the PE routers. Neither CE routers nor provider routers are required to perform any VPN functions.

MPLS VPN Routing

VPNs tunnel traffic as follows from one customer site to another customer site, using a public network as a transit network, when certain requirements are met:

1. Traffic is forwarded by standard IP forwarding from the CE routers to the PE routers.
2. The PE routers establish an LSP through the provider network.
3. The inbound PE router receives traffic, and it performs a route lookup. The lookup yields an LSP next hop, and the traffic is forwarded along the LSP.
4. The traffic reaches the outbound PE router, and the PE router pops the MPLS label and forwards the traffic with standard IP routing.

VRF Instances

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The interfaces belong to the routing tables, and the routing protocol parameters control the information in the routing tables. In the case of MPLS VPNs, each VPN has a VPN routing and forwarding (VRF) instance.

A VRF instance consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table. Because each instance is configured for a particular VPN, each VPN has separate tables, rules, and policies that control its operation.

A separate VRF table is created for each VPN that has a connection to a CE router. The VRF table is populated with routes received from directly connected CE sites associated with the VRF instance, and with routes received from other PE routers in the same VPN.

Route Distinguishers

Because a typical transit network is configured to handle more than one VPN, the provider routers are likely to have multiple VRF instances configured. As a result, depending on the origin of the traffic and any filtering rules applied to the traffic, the BGP routing tables can contain multiple routes for a particular destination address. Because BGP requires that exactly one BGP route per destination be imported into the forwarding table, BGP must have a way to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs.

A route distinguisher is a locally unique number that identifies all route information for a particular VPN. Unique numeric identifiers allow BGP to distinguish between routes that are otherwise identical.

Each routing instance that you configure on a PE router must have a unique route distinguisher. There are two possible formats:

- **as-number:number**, where **as-number** is an autonomous system (AS) number (a 2-byte value) in the range 1 through 65,535, and **number** is any 4-byte value. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the ISP or the customer AS number.
- **ip-address:number**, where **ip-address** is an IP address (a 4-byte value) and **number** is any 2-byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the **router-id** statement, which is a public IP address in your assigned prefix range.

The route target defines which route is part of a VPN. A unique route target helps distinguish between different VPN services on the same router. Each VPN also has a policy that defines how routes are imported into the VRF table on the router. A Layer 2 VPN is configured with import and export policies. A Layer 3 VPN uses a unique route target to distinguish between VPN routes.

The PE router then exports the route in IBGP sessions to the other provider routers. Route export is governed by any routing policy that has been applied to the particular VRF table. To propagate the routes through the provider network, the PE router must also convert the route to VPN format, which includes the route distinguisher.

When the outbound PE router receives the route, it strips off the route distinguisher and advertises the route to the connected CE router, typically through standard BGP IPv4 route advertisements.

MPLS Exceptions on SRX Series Devices

The MPLS implementation on SRX Series device is similar to MPLS implementations on M Series, T Series, and MX Series routers, with the following exception:

- SRX Series devices do not support aggregated Ethernet interfaces. Therefore, aggregated Ethernet interfaces between CE devices and PE routers are not supported for MPLS implementations of Layer 2 VPNs and Layer 2 Circuits on SRX Series devices.

RELATED DOCUMENTATION

[Understanding MPLS Layer 2 VPNs | 1427](#)

[Understanding MPLS Layer 3 VPNs | 1471](#)

[Understanding MPLS Layer 2 Circuits | 1463](#)

Understanding 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*.

Understanding Using MPLS-Based Layer 3 VPNs on Switches

IN THIS SECTION

- [MPLS-Based Layer 3 VPNs | 1420](#)

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

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 ["Understanding Interprovider and Carrier-of-Carriers VPNs" on page 1647](#)

RELATED DOCUMENTATION

[Understanding MPLS Label Operations | 422](#)

[Understanding MPLS Components for QFX Series and EX4600 Switches | 38](#)

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

[Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches | 1498](#)

Configuring a BGP Session for MPLS VPNs (CLI Procedure)

NOTE: This section is valid for Layer 2 VPNs and Layer 3 VPNs, but not Layer 2 circuits.

To configure an IBGP session, perform the following steps on each PE router:

1. Configure BGP.

[edit]


```
user@host# edit protocols bgp group group-name
```

2. Set the BGP type to internal.

```
[edit protocols bgp group group-name]  
user@host# set type internal
```

3. Specify the loopback interface.

```
[edit protocols bgp group group-name]  
user@host# set local-address loopback-interface-ip-address
```

4. Set the Layer 2 or Layer 3 VPN family type to unicast.

```
[edit protocols bgp group group-name]  
user@host# set family family-type unicast
```

Replace *family-type* with **l2vpn** for a Layer 2 VPN or **inet-vpn** for a Layer 3 VPN.

5. Enter the loopback address of the neighboring PE router.

```
[edit protocols bgp]  
user@host# set neighbor ip-address
```

6. Commit the configuration if you are finished configuring the device.

```
[edit]  
user@host# commit
```

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 3 VPN Configuration Overview | 1471](#)

Configuring an IGP and the RSVP Signaling Protocol (CLI Procedure)

To configure RSVP and OSPF:

1. Configure OSPF with traffic engineering support on the PE routers.

```
[edit]
user@host# edit protocols ospf traffic-engineering shortcuts
```

NOTE: You must configure the IGP at the **[edit protocols]** level, not within the routing instance at the **[edit routing-instances]** level.

2. Enable RSVP on interfaces that participate in the LSP. For PE routers, enable interfaces on the source and destination points. For provider routers, enable interfaces that connect the LSP between the PE routers.

```
[edit]
user@host# edit protocols rsvp interface interface-name
```

3. Commit the configuration if you are finished configuring the device.

```
[edit]
user@host# commit
```

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 3 VPN Configuration Overview | 1471](#)

[MPLS Layer 2 Circuit Configuration Overview | 1464](#)

Configuring Routing Options for MPLS VPNs (CLI Procedure)

To configure routing options for a VPN:

1. Configure the AS number.


```
[edit]
user@host# set routing-options autonomous-system as-number
```

2. Commit the configuration if you are finished configuring the device.

```
[edit]
user@host# commit
```

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 3 VPN Configuration Overview | 1471](#)

[MPLS Layer 2 Circuit Configuration Overview | 1464](#)

Configuring a Routing Instance for MPLS VPNs (CLI Procedure)

To configure a VPN routing instance on each PE router:

1. Create the routing instance.

```
[edit]
user@host# edit routing-instances routing-instance-name
```

2. Create a routing instance description. (This text appears in the output of the **show route instance detail** command.)

```
[edit routing-instances routing-instance-name]
user@host# set description "text"
```

3. Specify the instance type, either **l2vpn** for Layer 2 VPNs or **vrf** for Layer 3 VPNs.

```
[edit routing-instances routing-instance-name]
user@host# set instance-type instance-type
```

4. Specify the interface of the remote PE router.


```
[edit routing-instances routing-instance-name]  
user@host# set interface interface-name
```

5. Specify the route distinguisher using one of the following commands:

```
[edit routing-instances routing-instance-name]  
user@host# set route-distinguisher as-number:number  
user@host# set route-distinguisher ip-address:number
```

6. Specify the policy for the Layer 2 VRF table.

```
[edit routing-instances routing-instance-name]  
user@host# set vrf-import import-policy-name vrf-export export-policy-name
```

7. Specify the policy for the Layer 3 VRF table.

```
[edit routing-instances routing-instance-name]  
user@host# set vrf-target target:community-id
```

Where *community-id* is either *as-number:number* or *ip-address:number*.

8. Commit the configuration if you are finished configuring the device.

```
[edit]  
user@host# commit
```

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 3 VPN Configuration Overview | 1471](#)

Chained Composite Next Hops for Transit Devices for VPNs

The Juniper Networks PTX Series Packet Transport Routers, MX Series 5G Universal Routing Platforms 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 5G Universal Routing Platforms containing both DPC and MPC FPCs, chained composite next hops are disabled by default.

To enable chained composite next hops on the MX Series routers such as: MX80, MX240, MX480, MX960 and MX2020, the chassis must be configured to use the **enhanced-ip** option in network services mode.

On T4000 Core Routers containing MPC and FPCs, chained composite next hops are disabled by default.

To enable chained composite next hops on a T4000 router, the chassis must be configured to use the **enhanced-mode** option in network services mode.

For more information about configuring chassis network services, see the *Junos OS Administration Library*.

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 \(Chained Composite Next Hops\) | 3352](#)

Understanding MPLS Layer 2 VPNs

In an MPLS Layer 2 VPN, traffic is forwarded to the provider edge (PE) router in Layer 2 format, carried by MPLS through an label-switched path (LSP) over the service provider network, and then converted back to Layer 2 format at the receiving customer edge (CE) router.

Routing occurs on the customer routers, typically on the CE router. The CE router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The PE router receiving the traffic sends it across the network to the PE router on the outbound side. The PE routers need no information about the customer's routes or routing topology, and need only to determine the virtual tunnel through which to send the traffic.

Implementing a Layer 2 VPN on the router is similar to implementing a VPN using a Layer 2 technology such as Asynchronous Transfer Mode (ATM) or Frame Relay.

RELATED DOCUMENTATION

[MPLS VPN Overview | 1416](#)

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

Understanding Ethernet-over-MPLS (L2 Circuit)

IN THIS SECTION

- [Ethernet-over-MPLS in Data Centers | 1428](#)

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

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 \(L2 Circuit\) | 1433](#)

MPLS Layer 2 VPN Configuration Overview

To configure MPLS Layer 2 VPN functionality on a router running Junos OS, you must enable support on the provider edge (PE) router and configure the PE router to distribute routing information to other routers in the VPN, as explained in the following steps. However, because the tunnel information is maintained at both PE routers, neither the provider core routers nor the customer edge (CE) routers need to maintain any VPN information in their configuration databases.

To configure an MPLS Layer 2 VPN:

1. Determine all of the routers that you want to participate in the VPN, and then complete the initial configuration of their interfaces. See *Interfaces User Guide for Security Devices*.
2. For all of the routers in the VPN configuration, update the interface configurations to enable participation in the Layer 2 VPN. As part of the interface configuration, you must configure the MPLS address family for each interface that uses LDP or RSVP. See [“Configuring Interfaces for Layer 2 VPNs \(CLI Procedure\)” on page 1432](#).
3. For all of the routers in the VPN configuration, configure the appropriate protocols.
 - a. MPLS—For PE routers and provider routers, use MPLS to advertise the Layer 2 VPN interfaces that communicate with other PE routers and provider routers.
 - b. BGP and internal BGP (IBGP)—For PE routers, configure a BGP session to enable the routers to exchange information about routes originating and terminating in the VPN. (The PE routers use this information to determine which labels to use for traffic destined to the remote sites. The IBGP session for the VPN runs through the loopback address.) In addition, CE routers require a BGP connection to the PE routers. See [“Configuring a BGP Session for MPLS VPNs \(CLI Procedure\)” on page 1421](#).
 - c. IGP and a signaling protocol—For PE routers, configure a signaling protocol (either LDP or RSVP) to dynamically set up label-switched paths (LSPs) through the provider network. (LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed.) You must use LDP or RSVP between PE routers and provider routers, but you cannot use them for interfaces between PE routers and CE routers.

In addition, configure an IGP such as OSPF or static routes for PE routers to enable exchanges of routing information between the PE routers and provider routers. Each PE router's loopback address must appear as a separate route. Do not configure any summarization of the PE router's loopback addresses at the area boundary. Configure the provider network to run OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector.

See [“Configuring an IGP and the LDP Signaling Protocol \(CLI Procedure\)” on page 1466](#) and [“Configuring an IGP and the RSVP Signaling Protocol \(CLI Procedure\)” on page 1423](#).

4. For all of the routers in the VPN configuration, configure routing options. The only required routing option for VPNs is the AS number. You must specify it on each router involved in the VPN. See [“Configuring Routing Options for MPLS VPNs \(CLI Procedure\)” on page 1423](#).
5. For each PE router in the VPN configuration, configure a routing instance for each VPN. The routing instance should have the same name on each PE router. Each routing instance must have a unique route distinguisher associated with it. (VPN routing instances need a route distinguisher to help BGP distinguish between potentially identical network layer reachable information [NLRI] messages received from different VPNs.) See [“Configuring a Routing Instance for MPLS VPNs \(CLI Procedure\)” on page 1424](#).
6. For each PE router in the VPN configuration, configure a VPN routing policy if you are not using a route target. Within the policy, describe which packets are sent and received across the VPN and specify how routes are imported into and exported from the router's VRF table. Each advertisement must have an associated route target that uniquely identifies the VPN for which the advertisement is valid. If the routing instance uses a policy for accepting and rejecting packets instead of a route target, you must specify the import and export routing policies and the community on each PE router. See [“Configuring a Routing Policy for MPLS Layer 2 VPNs \(CLI Procedure\)” on page 1430](#).

RELATED DOCUMENTATION

[Verifying an MPLS Layer 2 VPN Configuration | 1459](#)

Configuring a Routing Policy for MPLS Layer 2 VPNs (CLI Procedure)

These instructions show how to configure a Layer 2 VPN routing policy on the PE routers in the VPN.

After configuring an import routing policy for a Layer 2 VPN, configure an export routing policy for the Layer 2 VPN. Configure this export policy on the PE routers in the VPN. The export routing policy defines how routes are exported from the PE router routing table. An export policy is applied to routes sent to other PE routers in the VPN. The export policy must also evaluate all routes received over the routing protocol session with the CE router. The export policy must also contain a second term for rejecting all other routes.

To configure a Layer 2 VPN routing policy on a PE router:

1. Configure the import routing policy.


```
[edit]
user@host# edit policy-options policy-statement import-policy-name
```

2. Define the import policy's term for accepting packets.

```
[edit edit policy-options policy-statement import-policy-name]
user@host# set term term-name-accept from protocol bgp community community-name
user@host# set term term-name-accept then accept
```

3. Define the import policy's term for rejecting packets.

```
[edit edit policy-options policy-statement import-policy-name]
user@host# set term term-name-reject then reject
```

4. Configure the export routing policy.

```
[edit]
user@host# edit policy-options policy-statement export-policy-name
```

5. Define the export policy's term for accepting packets.

```
[edit policy-options policy-statement export-policy-name]
user@host# set term term-name-accept from community add community-name
user@host# set term term-name-accept then accept
```

6. Define the export policy's term for rejecting packets.

```
[edit policy-options policy-statement export-policy-name]
user@host# set term term-name-reject from community add community-name
user@host# set term term-name-reject then reject
```

7. Define the export policy's community using one of the following commands.

```
[edit policy-options policy-statement export-policy-name]
user@host# community community-name target: as-number
user@host# community community-name target: ip-address:number
```


8. Commit the configuration if you are finished configuring the device.

```
[edit]
user@host# commit
```

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 3 VPN Configuration Overview | 1471](#)

Configuring Interfaces for Layer 2 VPNs (CLI Procedure)

Configuring the device interfaces that participate in the VPN is similar to configuring them for other uses, with a few requirements for the VPN. Perform the following tasks for each interface involved in the VPN, except Layer 3 loopback interfaces, which do not require other configuration.

To configure an interface for an MPLS VPN:

1. Configure IPv4 address and loopback address on all of the interfaces of the provider edge (PE) and provider (P) devices, excluding the Layer 2 VPN interfaces facing a customer edge (CE) device:

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family inet address ipv4_address
user@host# set lo0 unit logical-unit-number family inet address ipv4_address primary
```

2. Enable MPLS address family on all the PE and P device interfaces, excluding the loopback interfaces and the CE-facing Layer 2 VPN interfaces.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family mpls
```

3. For CE-facing Layer 2 VPN interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

```
[edit interfaces]
user@host# set interfaces interface-name vlan-tagging encapsulation vlan-ccc unit logical-unit-number
encapsulation vlan-ccc vlan-id vlan-id-number
```


4. Configure the physical and logical link-layer encapsulation type on all the PE devices that communicate with the CE devices on Layer 2 VPNs and Layer 2 circuits.

```
[edit interfaces]
user@host# set interface-name encapsulation encapsulation-type
user@host# set interface-name unit logical-unit-number encapsulation encapsulation-type
```

5. Commit the configuration if you are finished configuring the device.

```
[edit]
user@host# commit
```

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 3 VPN Configuration Overview | 1471](#)

[MPLS Layer 2 Circuit Configuration Overview | 1464](#)

Configuring Ethernet over MPLS (L2 Circuit)

IN THIS SECTION

- [Configuring the Local PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) | 1435](#)
- [Configuring the Remote PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) | 1436](#)
- [Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit | 1436](#)
- [Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit | 1437](#)


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 111 on page 1434 shows an example of a Layer 2 circuit configuration.

Figure 111: 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: 127.1.1.1
- PE2: 127.1.1.2

NOTE: On QFX Series and EX4600 switches, the L2 circuit CE facing interface does not support AE interfaces.

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 a 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
```

2. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch# set l2circuit neighbor 127.1.1.1 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 127.1.1.1
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls interface xe-0/0/7
user@switch# set 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
```

2. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch# set l2circuit neighbor 127.1.1.2 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 127.1.1.2
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls 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
```

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 127.1.1.1 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 127.1.1.1
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set 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
```

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 127.1.1.2 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 127.1.1.2
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls interface xe-0/0/7
user@switch# set ldp interface lo0.0
```

RELATED DOCUMENTATION

[Understanding Ethernet-over-MPLS \(L2 Circuit\) | 1428](#)

Example: Configuring MPLS-Based Layer 2 VPNs

IN THIS SECTION

- [Requirements | 1439](#)
- [Overview and Topology | 1439](#)
- [Configuring the Local PE Routing Device | 1443](#)
- [Configuring the Remote PE Routing Device | 1447](#)
- [Verification | 1450](#)

You can implement an MPLS-based Layer 2 virtual private network (VPN) using Junos OS routing devices to interconnect customer sites with Layer 2 technology. Layer 2 VPNs give customers complete control of their own routing. To support an MPLS-based Layer 2 VPN, you need to add components to the configuration of the two provider edge (PE) routing devices. You do not need to change the configuration of the provider devices.

This example shows how to configure an MPLS-based Layer 2 VPN.

NOTE: You can configure both an MPLS-based Layer 2 VPN and an MPLS-based Layer 3 VPN on the same device. However, you cannot configure the same customer edge interface to support both a Layer 2 VPN and a Layer 3 VPN. The core interfaces and the loopback interfaces are configured in the same way for Layer 2 VPNs and Layer 3 VPNs.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 11.1 or later if you are using EX Series switches
- Two PE routing devices

Before you configure the Layer 2 VPN components, configure the basic components for an MPLS network:

- Configure two PE routing devices. See [“Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 95](#).
- Configure one or more provider devices. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

NOTE: A Layer 2 VPN requires that the PE routing devices be configured using circuit cross-connect (CCC). The provider routing devices are configured in the same way for MPLS using CCC and for IP over MPLS.

Overview and Topology

A Layer 2 VPN provides complete separation between the provider’s network and the customer’s network—that is, the PE devices and the CE devices do not exchange routing information. Some benefits of a Layer 2 VPN are that it is private, secure, and flexible.

This example shows how to configure Layer 2 VPN components on the local and remote PE devices. This example does not include configuring a provider device, because there are no specific Layer 2 VPN components on the provider devices.

In the basic MPLS configuration of the PE devices using a circuit cross-connect (CCC), the PE devices are configured to use an interior gateway protocol (IGP), such as OSPF or IS-IS, as the routing protocol between the MPLS devices and LDP or RSVP as the signaling protocol. Traffic engineering is enabled. A label-switched path (LSP) is configured within the **[edit protocols]** hierarchy. However, unlike the basic MPLS configuration using a CCC, you do not need to associate the LSP with the customer edge interface. When you are configuring a Layer 2 VPN, you must use BGP signaling. The BGP signaling automates the connections, so manual configuration of the association between the LSP and the customer edge interface is not required.

The following components must be added to the PE routing devices for an MPLS-based Layer 2 VPN:

- BGP group with **family l2vpn signaling**
- Routing instance using instance type **l2vpn**
- The physical layer encapsulation type (**ethernet**) must be specified on the customer edge interface and the encapsulation type must also be specified in the configuration of the routing instance.

Figure 112 on page 1440 illustrates the topology of this MPLS-based Layer 2 VPN.

Figure 112: MPLS-Based Layer 2 VPN

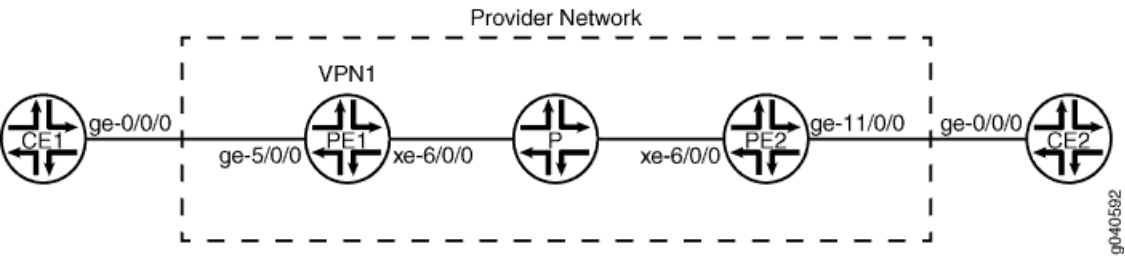


Table 33 on page 1440 shows the settings of the customer edge interface on the local CE device.

Table 33: Local CE Routing Device in the MPLS-Based Layer 2 VPN Topology

Property	Settings	Description
Local CE routing device hardware	Routing device	CE1
Customer edge interface	ge-0/0/0 unit 0 family inet address 10.0.0.2/16	Interface that connects CE1 to PE1.

Table 34 on page 1441 shows the settings of the customer edge interface on the remote CE routing device.

Table 34: Remote CE Routing Device in the MPLS-Based Layer 2 VPN Topology

Property	Settings	Description
Remote CE routing device hardware	Routing device	CE2
Customer edge interface	ge-0/0/0 unit 0 family inet address 10.0.0.1/16	Interface that connects CE2 to PE2.

[Table 35 on page 1441](#) shows the Layer 2 VPN components of the local PE routing device.

Table 35: Layer 2 VPN Components of the Local PE Routing Device

Property	Settings	Description
Local PE routing device hardware	Routing device	PE1
Customer edge interface	ge-5/0/0 encapsulation ethernet-ccc unit 0 family ccc	Connects PE1 to CE1. For the Layer 2 VPN, add ethernet-ccc as the physical layer encapsulation type. NOTE: The family ccc should already have been completed as part of the basic MPLS configuration of a PE routing device for circuit cross-connect. It is included here to show what was specified for that portion of the configuration.
Core interface	xe-6/0/0 unit 0 family inet address 10.0.0.60/16 family iso family mpls	Connects PE1 to P. NOTE: 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 family inet address 192.0.2.0/24 family iso address 49.0001.2102.2021.0210.00	NOTE: 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 35: Layer 2 VPN Components of the Local PE Routing Device (*continued*)

Property	Settings	Description
BGP	bgp	Added for the Layer 2 VPN configuration.
Routing instance	vpn1	Added for the Layer 2 VPN configuration

[Table 36 on page 1442](#) shows the Layer 2 VPN components of the remote PE routing device.

Table 36: Layer 2 VPN Components of the Remote PE Routing Device

Property	Settings	Description
PE routing device hardware	Routing device	PE2
Customer edge interface	ge-11/0/0 encapsulation ethernet-ccc unit 0 family ccc	Connects PE2 to CE2. For the Layer 2 VPN, add ethernet-ccc as the physical layer encapsulation type. NOTE: The family ccc should already have been completed as part of the basic MPLS configuration of a PE routing device for circuit cross-connect. It is included here to show what was specified for that portion of the configuration.
Core interface	xe-6/0/0 unit 0 family inet address 10.2.0.61/16 family iso family mpls	Connects PE2 to P. NOTE: 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 family inet address 192.0.2.3/24 family iso address 49.0001.2202.2022.0220.00	NOTE: 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 36: Layer 2 VPN Components of the Remote PE Routing Device (*continued*)

Property	Settings	Description
BGP	bgp	Added for the Layer 2 VPN configuration.
Routing instance	vpn1	Added for the Layer 2 VPN configuration.

Configuring the Local PE Routing Device

CLI Quick Configuration

To quickly configure the Layer 2 VPN components on the local PE routing device, copy the following commands and paste them into the routing device terminal window:

```
[edit]
set interfaces ge-5/0/0 encapsulation ethernet-ccc
set protocols bgp group ibgp local-address 192.0.2.0 family l2vpn signaling
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 192.0.2.3
set routing-instances vpn1 instance-type l2vpn
set routing-instances vpn1 interface ge-5/0/0
set routing-instances vpn1 route-distinguisher 192.0.2.0:21
set routing-instances vpn1 vrf-target target:21:21
set routing-instances vpn1 protocols l2vpn encapsulation-type ethernet
set routing-instances vpn1 protocols l2vpn interface ge-5/0/0.0 description "BETWEEN PE1 AND CE1"
set routing-instances vpn1 protocols l2vpn site JE-V21 site-identifier 21 interface ge-5/0/0 remote-site-id
26
```

Step-by-Step Procedure

To configure the Layer 2 VPN components on the local PE routing device:

1. Configure the customer edge interface to use the physical encapsulation type **ethernet-ccc**:

```
[edit]
user@PE1# set interfaces ge-5/0/0 encapsulation ethernet-ccc
```

2. Configure BGP, specifying the loopback address as the local address and enabling **family l2vpn signaling**:

```
[edit protocols bgp]
user@PE1# set group ibgp local-address 192.0.2.0 family l2vpn signaling
```


3. Configure the BGP group, specifying the group name and type:

```
[edit protocols bgp]
user@PE1# set group ibgp type internal
```

4. Configure the BGP neighbor, specifying the loopback address of the remote PE routing device as the neighbor's address:

```
[edit protocols bgp]
user@PE1# set group ibgp neighbor 192.0.2.3/24
```

5. Configure the routing instance, specifying the routing-instance name and using **l2vpn** as the instance type:

```
[edit routing-instances]
user@PE1# set vpn1 instance-type l2vpn
```

6. Configure the routing instance to apply to the customer edge interface:

```
[edit routing-instances]
user@PE1# set vpn1 interface ge-5/0/0
```

7. Configure the routing instance to use a route distinguisher:

```
[edit routing-instances]
user@PE1# set vpn1 route-distinguisher 192.0.2.0:21
```

8. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@PE1# set vpn1 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 Configuration Guide](#).

9. Configure the protocols and encapsulation type used by the routing instance:

```
[edit routing-instances]
user@PE1# set vpn1 protocols l2vpn encapsulation-type ethernet
```

10. Apply the routing instance to a customer edge interface and specify a description for it:


```
[edit routing-instances]
user@PE1# set vpn1 protocols interface ge-5/0/0.0 description "BETWEEN PE1 AND CE1"
```

11. Configure the routing-instance protocols site:

```
[edit routing-instances]
user@PE1# set vpn1 protocols l2vpn site JE-V21 site-identifier 21 remote-site-id 26
```

NOTE: The remote site ID (configured with the **remote-site-id** statement) corresponds to the site ID (configured with the **site-identifier** statement) configured on the other PE routing device.

Results

Display the results of the configuration:

```
user@PE1# show
```

```
interfaces {
  ge-5/0/0 {
    encapsulation ethernet-ccc;
    unit 0 {
      family ccc;
    }
  }
  xe-6/0/0 {
    unit 0 {
      family inet {
        address 10.0.0.60/16;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.0.2.0/24;
      }
      family iso {
        address 49.0001.2102.2021.0210.00;
      }
    }
  }
}
```



```

    }
}
protocols {
    rsvp {
        interface lo0.0;
        interface xe-0/0/6.0;
    }
    mpls {
        label-switched-path lsp_to_pe2 {
            to 192.0.2.3;
        }
    }
    bgp {
        group ibgp
        type internal
        local-address 192.0.2.0
        family l2vpn signaling
    }
}
routing-instances {
    vpn1 {
        instance-type l2vpn;
        interface ge-5/0/0.0;
        route-distinguisher 192.0.2.0:21;
        vrf-target target:21:21;
        protocols {
            l2vpn {
                encapsulation-type ethernet;
                interface ge-5/0/0.0 {
                    description "BETWEEN PE1 AND CE1";
                }
                site JE-V21 {
                    site-identifier 21;
                    interface ge-5/0/0.0 {
                        remote-site-id 26;
                    }
                }
            }
        }
    }
}
}
}
}

```


Configuring the Remote PE Routing Device

CLI Quick Configuration

To quickly configure the Layer 2 VPN components on the remote PE routing device, copy the following commands and paste them into the routing device terminal window:

```
[edit]
set interfaces ge-11/0/0 encapsulation ethernet-ccc
set protocols bgp group ibgp local-address 192.0.2.3 family l2vpn signaling
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 192.0.2.0
set routing-instances vpn1 instance-type l2vpn
set routing-instances vpn1 interface ge-11/0/0
set routing-instances vpn1 route-distinguisher 192.0.2.0:21
set routing-instances vpn1 vrf-target target:21:21
set routing-instances vpn1 protocols l2vpn encapsulation-type ethernet
set routing-instances vpn1 protocols l2vpn interface ge-11/0/0.0 description "BETWEEN PE1 AND CE1"
set routing-instances vpn1 protocols l2vpn site T26-VPN1 site-identifier 26 remote-site-id 21
```

Step-by-Step Procedure

To configure the Layer 2 VPN components on the remote PE routing device:

1. Configure the customer edge interface to use the physical encapsulation type **ethernet-ccc**:

```
[edit]
user@PE1# set interfaces ge-11/0/0 encapsulation ethernet-ccc
```

2. Configure BGP, specifying the loopback address as the **local-address** and specifying **family l2vpn signaling**:

```
[edit protocols bgp]
user@PE2# set group ibgp local-address 192.0.2.3 family l2vpn signaling
```

3. Configure the BGP group, specifying the group name and type:

```
[edit protocols bgp]
user@PE2# set group ibgp type internal
```

4. Configure the BGP neighbor, specifying the loopback address of the remote PE routing device as the neighbor's address:

```
[edit protocols bgp]
user@PE2# set group ibgp neighbor 192.0.2.0
```


5. Configure the routing instance, specifying the routing-instance name and using **l2vpn** as the **instance-type**:

```
[edit routing-instances]
user@PE2# set vpn1 instance-type l2vpn
```

6. Configure the routing instance to apply to the customer edge interface:

```
[edit routing-instances]
user@PE2# set vpn1 interface ge-11/0/0.0
```

7. Configure the routing instance to use a route distinguisher, using the format *ip-address:number*:

```
[edit routing-instances]
user@PE2# set vpn1 route-distinguisher 192.0.2.0:21
```

8. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@PE2# set vpn1 vrf-target target:21:21
```

9. Configure the protocols and encapsulation type used by the routing instance:

```
[edit routing-instances]
user@PE2# set vpn1 protocols l2vpn encapsulation-type ethernet
```

10. Apply the routing instance to a customer edge interface and specify a description for it:

```
[edit routing-instances]
user@PE1# set vpn1 protocols interface ge-11/0/0.0 description "BETWEEN PE1 AND CE1"
```

11. Configure the routing-instance protocols site:

```
[edit routing-instances]
user@PE2# set vpn1 protocols l2vpn site T26-VPN1 site-identifier 26 remote-site-id 21
```

NOTE: The remote site ID (configured with the **remote-site-id** statement) corresponds to the site ID (configured with the **site-identifier** statement) configured on the other PE routing device.

Results

Display the results of the configuration:

user@PE2# **show**

```
interfaces {
  ge-11/0/0 {
    encapsulation ethernet-ccc;
    unit 0 {
      family ccc;
    }
  }
  xe-6/0/0 {
    unit 0 {
      family inet {
        address 10.2.0.61/16;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.0.2.3/24;
      }
      family iso {
        address 49.0001.2202.2022.0220.00;
      }
    }
  }
}
protocols {
  rsvp {
    interface lo0.0;
    interface xe-0/0/6.0;
  }
  mpls {
    label-switched-path lsp_to_pe1 {
      to 192.0.2.0;
    }
  }
  bgp {
    group ibgp
    type internal
    local-address 192.0.2.0
  }
}
```



```

        family l2vpn signaling
    }
}
routing-instances {
    vpn1 {
        instance-type l2vpn;
        interface ge-11/0/0.0;
        route-distinguisher 192.0.2.0:21;
        vrf-target target:21:21;
        protocols {
            l2vpn {
                encapsulation-type ethernet;
                interface ge-11/0/0.0 {
                    description "BETWEEN PE1 AND CE1";
                }
                site T26-VPN1 {
                    site-identifier 26;
                    interface ge-11/0/0.0 {
                        remote-site-id 21;
                    }
                }
            }
        }
    }
}
}

```

Verification

IN THIS SECTION

- [Verifying the Layer 2 VPN Connection | 1451](#)
- [Verifying the Status of MPLS Label-Switched Paths | 1452](#)
- [Verifying BGP Status | 1452](#)
- [Verifying the Status of the RSVP Sessions | 1453](#)
- [Verifying the Routes in the Routing Table | 1454](#)
- [Pinging the Layer 2 VPN Connections | 1455](#)

To confirm that the MPLS-based Layer 2 VPN is working properly, perform these tasks:

Verifying the Layer 2 VPN Connection

Purpose

Verify that the Layer 2 VPN connection is up.

Action

user@PE1> show l2vpn connections

```

Layer-2 VPN connections:

Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned    <- -- only inbound connection is up
OR -- out of range              Up -- operational
OL -- no outgoing label         Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch              MI -- Mesh-Group ID not available
BK -- Backup connection         ST -- Standby connection
PF -- Profile parse failure     PB -- Profile busy
RS -- remote site standby       SN -- Static Neighbor

Legend for interface status
Up -- operational
Dn -- down

Instance: vpn1
  Local site: JE-V21 (21)
    connection-site      Type  St      Time last up      # Up trans
    26                   rmt   Up      Apr 16 05:53:21 2010      1
    Remote PE: 192.0.2.3, Negotiated control-word: Yes (Null)
    Incoming label: 800000, Outgoing label: 800001
    Local interface: ge-5/0/0.0, Status: Up, Encapsulation: ETHERNET

```

Meaning

The **St** field in the output shows that the Layer 2 VPN connection to **Remote PE (192.0.2.3)** is up.

Verifying the Status of MPLS Label-Switched Paths

Purpose

Verify that the MPLS label-switched paths (ingress and egress) are up.

Action

user@PE1> **show mpls lsp**

Ingress LSP: 1 sessions

To	From	State	Rt	P	ActivePath	LSPname
192.0.2.3	192.0.2.0	Up	0	*		lsp_to_pe2

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
192.0.2.0	192.0.2.3	Up	0	1 FF	3	-	lsp_to_pe1

Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Meaning

The **State** field in the output shows that the Ingress LSP to **Remote PE (192.0.2.3)** is up, and the Egress LSP from the remote PE routing device to this PE routing device (**192.0.2.0**) is also up.

Verifying BGP Status

Purpose

Verify that BGP is up.

Action

user@PE1> **show bgp summary**


```

Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l2vpn.0          1          1          0          0          0          0
Peer              AS      InPkt    OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
192.0.2.3           10         33         34         0         1        13:24 Establ

    bgp.l2vpn.0: 1/1/1/0
    vpn2.l2vpn.0: 1/1/1/0

```

Meaning

The output shows that the remote PE routing device (**192.0.2.3**) is listed as the BGP peer and that a protocol session has been established. It also shows the number of packets received from the remote PE routing device (**33**) and the number of packets sent (**34**) to the remote PE routing device.

Verifying the Status of the RSVP Sessions

Purpose

Verify that the RSVP sessions (ingress and egress) are up.

Action

```
user@PE1> show rsvp session
```

```

Ingress RSVP: 1 sessions
To          From          State    Rt  Style  Labelin Labelout LSPname
192.0.2.3    192.0.2.0    Up        0  1 FF      -    462880 lsp_to_pe2
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To          From          State    Rt  Style  Labelin Labelout LSPname
192.0.2.0    192.0.2.3    Up        0  1 FF      3      -    lsp_to_pe1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The output shows that both the ingress RSVP session and the egress RSVP session are up.

Verifying the Routes in the Routing Table

Purpose

On routing device PE 1, use the **show route table** command to verify that the routing table is populated with the Layer 2 VPN routes used to forward the traffic.

Action

user@PE1> **show route table** bgp.l2vpn.0

```
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192:2:27:27/96
    *[BGP/170] 00:13:55, localpref 100, from 192.0.2.3
    AS path: I
    > to 10.2.0.24 via ge-6/0/46.0, label-switched-path lsp_to_pe2
```

user@PE1> **show route table** vpn1.l2vpn.0

```
vpn1.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192:2:27:27/96
    *[BGP/170] 00:14:00, localpref 100, from 192.0.2.3
    AS path: I
    > to 10.2.0.24 via ge-6/0/46.0, label-switched-path lsp_to_pe2

192:2:28:27/96
    *[L2VPN/170/~101] 00:15:55, metric2 1
    Indirect
```

Meaning

The command **show route table bgp.l2vpn.0** displays all Layer 2 VPN routes that have been created on this routing device. The command **show route table vpn1.l2vpn.0** shows the Layer 2 VPN routes that have been created for the routing instance **vpn1**.

Pinging the Layer 2 VPN Connections

Purpose

Verify connectivity.

Action

```
user@PE1> ping mpls l2vpn interface xe-6/0/0.0 reply-mode ip-udp
```

```
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

```
user@PE1> ping mpls l2vpn instance vpn1 remote-site-id 26 local-site-id 21 detail
```

```
Request for seq 1, to interface 68, labels <800001, 100176>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <800001, 100176>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <800001, 100176>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <800001, 100176>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <800001, 100176>
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Meaning

The output shows that connectivity is established.

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches | 1498](#)

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches

IN THIS SECTION

- [MPLS-Based Layer 2 VPNs | 1456](#)
- [Layer 2 Circuits | 1457](#)
- [MPLS-Based Layer 3 VPNs | 1457](#)
- [Comparing an MPLS-Based Layer 2 VPN and an MPLS-Based Layer 3 VPN | 1458](#)

On EX8200 and EX4500 switches, you can use MPLS-based Layer 2 and Layer 3 virtual private networks (VPNs) or MPLS Layer 2 circuits, allowing you 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 2 VPNs

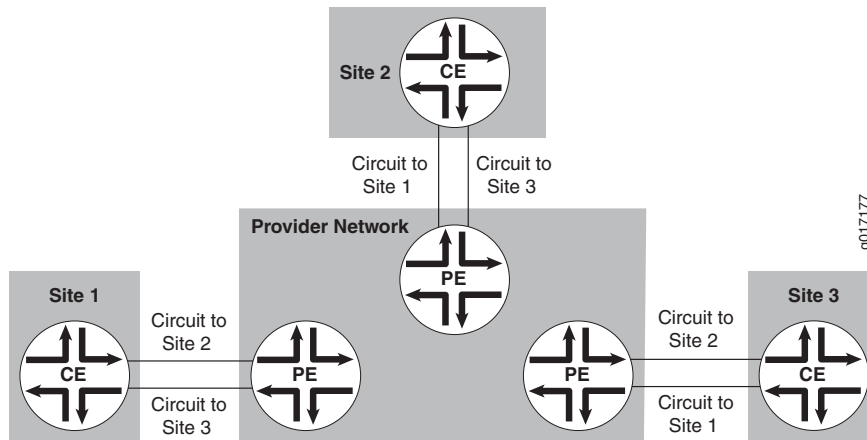
In an MPLS-based Layer 2 VPN, traffic is forwarded by the customer's customer edge (CE) switch (or router) to the service provider's provider edge (PE) switch in a Layer 2 format. It is carried by MPLS over the service provider's network and then converted back to Layer 2 format at the receiving site.

On a Layer 2 VPN, routing occurs on the customer's switches, typically on the CE switch. The CE switch connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The PE switch receiving the traffic sends it across the service provider's network to the PE switch connected to the receiving site. The PE switches do not store or process the customer's routes; the switches must be configured to send data to the appropriate tunnel.

For a Layer 2 VPN, customers must configure their own switches to carry all Layer 3 traffic. The service provider must detect only how much traffic the Layer 2 VPN will need to carry. The service provider's switches carry traffic between the customer's sites using Layer 2 VPN interfaces. The VPN topology is determined by policies configured on the PE switches.

Customers must know only which VPN interfaces connect to which of their own sites. [Figure 113 on page 1457](#) illustrates a full-mesh Layer 2 VPN in which each site has a VPN interface linked to each of the other customer sites. In a full-mesh topology between all three sites, each site requires two logical interfaces (one for each of the other CE routers or switches), although only one physical link is needed to connect each PE switch to each CE router or switch.

Figure 113: Layer 2 VPN Connecting CE Switches



Layer 2 Circuits

A Layer 2 circuit is a point-to-point Layer 2 connection that uses MPLS or another tunneling technology on the service provider's network. 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) switches. In contrast, each CCC requires a dedicated LSP.

The Junos OS implementation of Layer 2 circuits supports only the remote form of a Layer 2 circuit; that is, a connection from a local customer edge (CE) switch to a remote CE switch.

Packets are sent to the remote CE switch by means of an egress virtual private network (VPN) label advertised by the remote PE switch. The VPN label transits over either an RSVP or an LDP LSP (or other type) tunnel to the remote PE switch connected to the remote CE switch. LDP is the signaling protocol used for advertising VPN labels.

Return traffic sent from the remote CE switch to the local CE switch uses an ingress VPN label advertised by the local PE switch.

MPLS-Based Layer 3 VPNs

In a Layer 3 VPN, the routing occurs on the service provider's routers. Therefore, Layer 3 VPNs require more configuration on the part of the service provider, because the service provider's PE routers must store and process the customer's routes.

In the Junos OS, Layer 3 VPNs are based on RFC 4364, [BGP/MPLS IP Virtual Private Networks](#). This RFC defines a mechanism by which service providers can use their IP backbones to provide Layer 3 VPN services to their customers. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

VPNs based on RFC 4364 are also known as BGP/MPLS VPNs because BGP is used to distribute VPN routing information across the provider's backbone, and MPLS is used to forward VPN traffic across the backbone to remote VPN sites.

Customer networks, because they are private, can use either public addresses 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 private addresses used by other network users. BGP/MPLS VPNs solve this problem by prefixing a VPN identifier to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and within 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.

Comparing an MPLS-Based Layer 2 VPN and an MPLS-Based Layer 3 VPN

The differences between Layer 2 VPNs and Layer 3 VPNs are summarized in [Table 37 on page 1458](#)

Table 37: Comparing an MPLS-Based Layer 2 VPN and an MPLS-Based Layer 3 VPN

Layer 2 VPN	Layer 3 VPN
Customer sites appear to be on the same LAN even if geographically dispersed.	Service provider's technical expertise ensures efficient site-to-site routing.
	Service providers can provide additional value-added services through network convergence that encompasses voice, video, and data.

Table 37: Comparing an MPLS-Based Layer 2 VPN and an MPLS-Based Layer 3 VPN (*continued*)

Layer 2 VPN	Layer 3 VPN
<p>The service provider does not require information about the customer's network topology, policies, routing information, etc.</p> <p>The customer has complete control over policies and routing.</p>	<p>Customers must share information about their network topology.</p> <p>The service provider determines the policies and routing.</p>
<p>The CE switch forwards traffic to the service provider's PE switch in Layer 2 format.</p>	<p>The customer's CE switch must be configured to use BGP or OSPF to communicate with the service provider's PE switch to carry IP prefixes across the network. Other protocol packets are not supported.</p>

RELATED DOCUMENTATION

[Understanding MPLS Label Operations on EX Series Switches | 426](#)

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

[Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches | 1498](#)

Verifying an MPLS Layer 2 VPN Configuration

Purpose

Verify the connectivity of MPLS Layer 2 VPNs using the **ping mpls** command. This command helps to verify that a VPN has been enabled by testing the integrity of the VPN connection between the PE routers. It does not test the connection between a PE router and CE router.

Action

- To ping an interface configured for the Layer 2 VPN on the PE router, use the following command:

```
ping mpls l2vpn interface interface-name
```

- To ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier to test the integrity of the Layer 2 VPN connection (specified by identifiers) between the two PE routers, use the following command:

```
ping mpls l2vpn instance l2vpn-instance-name local-site-id local-site-id-number remote-site-id remote-site-id-number
```


RELATED DOCUMENTATION

| [MPLS Layer 2 VPN Configuration Overview](#) | 1429

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

You can configure MPLS-based Layer 2 virtual private networks (VPNs) on EX8200 and EX4500 switches. Some benefits of a Layer 2 VPN are that it is private, secure and flexible. To configure Layer 2 VPN functionality in your MPLS network, you must configure Layer 2 VPN components on the local and remote provider edge (PE) switches.

NOTE: This topic shows how to add Layer 2 VPN components to a CCC configured on a simple interface. For information on combining Layer 2 VPN components with a tagged VLAN CCC, see [“Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\)”](#) on page 1684.

Before you configure the Layer 2 VPN components, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\)”](#) on page 95.
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)”](#) on page 99.

NOTE: A Layer 2 VPN requires that the PE switches be configured using a circuit cross-connect (CCC).

Configure the Layer 2 VPN components on both PE switches. This procedure describes how to configure one PE switch. Repeat the procedure to configure the remote PE switch.

To configure Layer 2 VPN components on the PE switch:

1. Configure the customer edge interface to use the physical encapsulation type **ethernet-ccc**:

NOTE: The customer edge interface is a simple interface.

```
[edit]
user@switch# set interfaces interface-name encapsulation ethernet-ccc
```

2. Configure BGP, specifying the loopback address of this PE switch as the local address and specifying **family l2vpn signaling**:

```
[edit protocols bgp]
user@switch# set local-address address family l2vpn signaling
```

3. Configure the BGP group, specifying the group name and **type internal**:

```
[edit protocols bgp]
user@switch# set group group-name type internal
```

4. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:

```
[edit protocols bgp]
user@switch# set neighbor address
```

5. Configure the routing instance, specifying the routing-instance name and using **l2vpn** as the instance type:

```
[edit routing-instances]
user@switch# set routing-instance-name instance-type l2vpn
```

6. Configure the routing instance to apply to the customer edge interface:

```
user@switch# set routing-instances routing-instance-name interface interface-name
```

7. Configure the routing instance to use a route distinguisher:

NOTE: Each routing instance that you configure on a PE switch must have a unique route distinguisher associated with it. VPN routing instances must have 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.

```
user@switch# set routing-instances routing-instance-name route-distinguisher ip-address:number
```

8. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@switch# set routing-instance-name vrf-target community
```

NOTE: If you configure the **community** option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. 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 Configuration Guide](#).

9. Configure the protocols and encapsulation type used by the routing instance:

```
[edit routing-instances]
user@switch# set routing-instance-name protocols l2vpn encapsulation-type ethernet
```

10. Apply the routing instance to a customer edge interface and specify a description for it:

```
[edit routing-instances]
user@switch# set routing-instance-name protocols interface interface-name description text
```

11. Configure the routing instance protocols site:

```
[edit routing-instances]
user@switch# set routing-instance-name protocols l2vpn site site-name site-identifier
identifier remote-site-id identifier
```


NOTE: The remote site ID (configured with the **remote-site-id** statement) corresponds to the site ID (configured with the **site-identifier** statement) configured on the other PE switch.

RELATED DOCUMENTATION

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

[Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) | 1512](#)

[Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches | 1456](#)

Understanding MPLS Layer 2 Circuits

An MPLS Layer 2 circuit is a point-to-point Layer 2 connection that transports traffic by means of MPLS or another tunneling technology on the service provider network. The Layer 2 circuit creates a virtual connection to direct traffic between two customer edge (CE) routers across a service provider network. The main difference between a Layer 2 VPN and a Layer 2 circuit is the method of setting up the virtual connection. As with a leased line, a Layer 2 circuit forwards all packets received from the local interface to the remote interface.

Each Layer 2 circuit is represented by the logical interface connecting the local provider edge (PE) router to the local CE router. All Layer 2 circuits using a particular remote PE router neighbor is identified by its IP address and is usually the endpoint destination for the label-switched path (LSP) tunnel transporting the Layer 2 circuit.

Each virtual circuit ID uniquely identifies the Layer 2 circuit among all the Layer 2 circuits to a specific neighbor. The key to identifying a particular Layer 2 circuit on a PE router is the neighbor address and the virtual circuit ID. Based on the virtual circuit ID and the neighbor relationship, an LDP label is bound to an LDP circuit. LDP uses the binding for sending traffic on that Layer 2 circuit to the remote CE router.

RELATED DOCUMENTATION

[MPLS VPN Overview | 1416](#)

[MPLS Layer 2 Circuit Configuration Overview | 1464](#)

MPLS Layer 2 Circuit Configuration Overview

To configure an MPLS Layer 2 circuit:

1. Determine all of the routers that you want to participate in the circuit, and then complete the initial configuration of their interfaces. See the *Interfaces User Guide for Security Devices*.
2. For all of the routers in the circuit configuration, update the interface configurations to enable participation in the Layer 2 circuit.
 - a. On the interface communicating with the other provider edge (PE) router, specify MPLS and IPv4, and include the IP address. For the loopback interface, specify **inet**, and include the IP address. For IPv4, designate the loopback interface as primary so it can receive control packets. (Because it is always operational, the loopback interface is best able to perform the control function.)
 - b. On the PE router interface facing the customer edge (CE) router, specify a circuit cross-connect (CCC) encapsulation type. The type of encapsulation depends on the interface type. For example, an Ethernet interface uses **ethernet-ccc**. (The encapsulation type determines how the packet is constructed for that interface.)
 - c. On the CE router interface that faces the PE router, specify **inet** (for IPv4), and include the IP address. In addition, specify a routing protocol such as Open Shortest Path First (OSPF), which specifies the area and IP address of the router interface.

See [“Configuring Interfaces for Layer 2 VPNs \(CLI Procedure\)” on page 1432](#).

3. For all of the routers in the circuit configuration, configure the appropriate protocols.
 - a. MPLS—For PE routers and provider routers, use MPLS to advertise the Layer 2 circuit interfaces that communicate with other PE routers and provider routers.
 - b. BGP—For PE routers, configure a BGP session.
 - c. IGP and a signaling protocol—For PE routers, configure a signaling protocol (either LDP or RSVP) to dynamically set up label-switched paths (LSPs) through the provider network. (LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed.) You must use LDP or RSVP between PE routers and provider routers, but cannot use them for interfaces between PE routers and CE routers.

In addition, configure an IGP such as OSPF or static routes on the PE routers to enable exchanges of routing information between the PE routers and provider routers. Each PE router's loopback address must appear as a separate route. Do not configure any summarization of the PE router's

loopback addresses at the area boundary. Configure the provider network to run OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector.

See [“Configuring an IGP and the LDP Signaling Protocol \(CLI Procedure\)” on page 1466](#) and [“Configuring an IGP and the RSVP Signaling Protocol \(CLI Procedure\)” on page 1423](#).

4. For all of the routers in the circuit configuration, configure routing options. The only required routing option for circuits is the autonomous system (AS) number. You must specify it on each router involved in the circuit. See [“Configuring Routing Options for MPLS VPNs \(CLI Procedure\)” on page 1423](#).
5. For PE routers, configure Layer 2 circuits on the appropriate interfaces. See [“Configuring an MPLS Layer 2 Circuit \(CLI Procedure\)” on page 1465](#).

RELATED DOCUMENTATION

[Verifying an MPLS Layer 2 Circuit Configuration | 1466](#)

Configuring an MPLS Layer 2 Circuit (CLI Procedure)

To configure a Layer 2 circuit on a PE router:

1. Enable a Layer 2 circuit on the appropriate interface.

```
[edit]
user@host# edit protocols l2circuit neighbor interface-name interface interface-name
```

2. Enter the circuit ID number.

```
[edit protocols l2circuit neighbor interface-name interface interface-name]
user@host# set virtual-circuit-id id-number
```

For **neighbor**, specify the local loopback address, and for **interface**, specify the interface name of the remote PE router.

3. Commit the configuration if you are finished configuring the device.

```
[edit]
user@host# commit
```


RELATED DOCUMENTATION

[MPLS Layer 2 Circuit Configuration Overview | 1464](#)

Verifying an MPLS Layer 2 Circuit Configuration

Purpose

To verify the connectivity of MPLS Layer 2 circuits, use the **ping mpls** command. This command helps to verify that the circuit has been enabled by testing the integrity of the Layer 2 circuit between the source and destination routers.

Action

- To ping an interface configured for the Layer 2 circuit on the PE router, enter the following command:

```
ping mpls l2circuit interface interface-name
```

- To ping a combination of the IPv4 prefix and the virtual circuit ID on the destination PE router, enter the following command:

```
ping mpls l2circuit virtual-circuitprefix virtual-circuit-id
```

RELATED DOCUMENTATION

[MPLS Layer 2 Circuit Configuration Overview | 1464](#)

Configuring an IGP and the LDP Signaling Protocol (CLI Procedure)

The following instructions show how to configure LDP and OSPF on PE routers and provider routers. Within the task, you specify which interfaces to enable for LDP. Perform this step on each PE router interface and provider router interface that communicates with other PE routers and provider routers. For OSPF, you configure at least one area on at least one of the router's interfaces. (An AS can be divided into multiple areas.) These instructions use the backbone area **0.0.0.0** and show how to enable traffic engineering for Layer 2 VPN circuits.

To configure LDP and OSPF:

1. Enable the ldp protocol.


```
[edit]  
user@host# edit protocols ldp
```

NOTE: You must configure the IGP at the **[protocols]** level of the configuration hierarchy, not within the routing instance at the **[routing-instances]** level of the configuration hierarchy.

2. Specify which interfaces to enable for LDP.

```
[edit protocols ldp]  
user@host# edit interface interface-name
```

3. Configure OSPF for each interface that uses LDP.

```
[edit]  
user@host# edit protocols ospf area 0.0.0.0 interface interface-name
```

4. (Layer 2 VPN circuits only) Enable traffic engineering.

```
[edit protocols ospf]  
user@host# set traffic engineering
```

5. Commit the configuration if you are finished configuring the device.

```
[edit]  
user@host# commit
```

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 3 VPN Configuration Overview | 1471](#)

[MPLS Layer 2 Circuit Configuration Overview | 1464](#)

Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)

You can configure MPLS-based Layer 2 virtual private networks (VPNs) on EX8200 and EX4500 switches. Some benefits of a Layer 2 VPN are that it is private, secure and flexible. To configure Layer 2 VPN functionality in your MPLS network, you must configure Layer 2 VPN components on the local and remote provider edge (PE) switches.

NOTE: This topic shows how to add Layer 2 VPN components to a CCC configured on a simple interface. For information on combining Layer 2 VPN components with a tagged VLAN CCC, see [“Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\)” on page 1684](#).

Before you configure the Layer 2 VPN components, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 95](#).
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

NOTE: A Layer 2 VPN requires that the PE switches be configured using a circuit cross-connect (CCC).

Configure the Layer 2 VPN components on both PE switches. This procedure describes how to configure one PE switch. Repeat the procedure to configure the remote PE switch.

To configure Layer 2 VPN components on the PE switch:

1. Configure the customer edge interface to use the physical encapsulation type **ethernet-ccc**:

NOTE: The customer edge interface is a simple interface.

[edit]

```
user@switch# set interfaces interface-name encapsulation ethernet-ccc
```

2. Configure BGP, specifying the loopback address of this PE switch as the local address and specifying family **l2vpn** signaling:


```
[edit protocols bgp]
user@switch# set local-address address family l2vpn signaling
```

3. Configure the BGP group, specifying the group name and **type internal**:

```
[edit protocols bgp]
user@switch# set group group-name type internal
```

4. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:

```
[edit protocols bgp]
user@switch# set neighbor address
```

5. Configure the routing instance, specifying the routing-instance name and using **l2vpn** as the instance type:

```
[edit routing-instances]
user@switch# set routing-instance-name instance-type l2vpn
```

6. Configure the routing instance to apply to the customer edge interface:

```
user@switch# set routing-instances routing-instance-name interface interface-name
```

7. Configure the routing instance to use a route distinguisher:

NOTE: Each routing instance that you configure on a PE switch must have a unique route distinguisher associated with it. VPN routing instances must have 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.

```
user@switch# set routing-instances routing-instance-name route-distinguisher ip-address:number
```

8. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@switch# set routing-instance-name vrf-target community
```


NOTE: If you configure the **community** option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. 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 Configuration Guide](#).

9. Configure the protocols and encapsulation type used by the routing instance:

```
[edit routing-instances]
user@switch# set routing-instance-name protocols l2vpn encapsulation-type ethernet
```

10. Apply the routing instance to a customer edge interface and specify a description for it:

```
[edit routing-instances]
user@switch# set routing-instance-name protocols interface interface-name description text
```

11. Configure the routing instance protocols site:

```
[edit routing-instances]
user@switch# set routing-instance-name protocols l2vpn site site-name site-identifier
identifier remote-site-id identifier
```

NOTE: The remote site ID (configured with the **remote-site-id** statement) corresponds to the site ID (configured with the **site-identifier** statement) configured on the other PE switch.

RELATED DOCUMENTATION

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

[Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) | 1512](#)

[Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches | 1456](#)

Understanding MPLS Layer 3 VPNs

An MPLS Layer 3 VPN operates at the Layer 3 level of the OSI model, the Network layer. The VPN is composed of a set of sites that are connected over a service provider's existing public Internet backbone. The sites share common routing information and the connectivity of the sites is controlled by a collection of policies.

In an MPLS Layer 3 VPN, routing occurs on the service provider's routers. The provider routers route and forward VPN traffic at the entry and exit points of the transit network. The service provider network must learn the IP addresses of devices sending traffic across the VPN and the routes must be advertised and filtered throughout the provider network. As a result, Layer 3 VPNs require information about customer routes and a more extensive VPN routing and forwarding (VRF) policy configuration than a Layer 2 VPN. This information is used to share and filter routes that originate or terminate in the VPN.

The MPLS Layer 3 VPN requires more processing power on the provider edge (PE) routers than a Layer 2 VPN, because the Layer 3 VPN has larger routing tables for managing network traffic on the customer sites. Route advertisements originate at the customer edge (PE) routers and are shared with the inbound PE routers through standard IP routing protocols, typically BGP. Based on the source address, the PE router filters route advertisements and imports them into the appropriate VRF table.

The provider router uses OSPF and LDP to communicate with the PE routers. For OSPF, the provider router interfaces that communicate with the PE routers are specified, as well as the loopback interface. For the PE routers, the loopback interface is in passive mode, meaning it does not send OSPF packets to perform the control function.

RELATED DOCUMENTATION

[MPLS VPN Overview](#) | 1416

MPLS Layer 3 VPN Configuration Overview

To configure MPLS Layer 3 VPN functionality on a router running Junos OS, you must enable support on the provider edge (PE) router and configure the PE router to distribute routing information to other routers in the VPN, as explained in the following steps. However, because the tunnel information is maintained at both PE routers, neither the provider core routers nor the customer edge (CE) routers need to maintain any VPN information in their configuration databases.

To configure an MPLS Layer 3 VPN:

1. Determine all of the routers that you want to participate in the VPN, and then complete the initial configuration of their interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.
2. For all of the routers in the VPN configuration, update the interface configurations to enable participation in the Layer 3 VPN. As part of the interface configuration, you must configure the MPLS address family for each interface that uses LDP or RSVP. See [“Configuring Interfaces for Layer 2 VPNs \(CLI Procedure\)” on page 1432](#).
3. For all of the routers in the VPN configuration, configure the appropriate protocols.
 - a. MPLS—If you are using RSVP, use MPLS to advertise the Layer 3 VPN interfaces on the PE routers and provider routers that communicate with other PE routers and provider routers.
 - b. BGP, EBGp, and internal BGP (IBGP)—For PE routers, configure a BGP session to enable the routers to exchange information about routes originating and terminating in the VPN. (The PE routers use this information to determine which labels to use for traffic destined to the remote sites. The IBGP session for the VPN runs through the loopback address.) In addition, CE routers require a BGP connection to the PE routers. See [“Configuring a BGP Session for MPLS VPNs \(CLI Procedure\)” on page 1421](#).
 - c. IGP and a signaling protocol—For PE routers and provider, configure a signaling protocol (either LDP or RSVP) to dynamically set up label-switched paths (LSPs) through the provider network. (LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed.) You must use LDP or RSVP between PE routers and provider routers, but cannot use them for interfaces between PE routers and CE routers.

In addition, configure an IGP such as OSPF or static routes on the PE routers in order to enable exchanges of routing information between the PE routers and provider routers. Each PE router's loopback address must appear as a separate route. Do not configure any summarization of the PE router's loopback addresses at the area boundary. Configure the provider network to run OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector.

See [“Configuring an IGP and the LDP Signaling Protocol \(CLI Procedure\)” on page 1466](#) and [“Configuring an IGP and the RSVP Signaling Protocol \(CLI Procedure\)” on page 1423](#).
4. For all of the routers in the VPN configuration, configure routing options. The only required routing option for VPNs is the autonomous system (AS) number. You must specify it on each router involved in the VPN. See [“Configuring Routing Options for MPLS VPNs \(CLI Procedure\)” on page 1423](#).
5. For each PE router in the VPN configuration, configure a routing instance for each VPN. The routing instance should have the same name on each PE router. Each routing instance must have a unique route distinguisher associated with it. (VPN routing instances need a route distinguisher to help BGP

distinguish between potentially identical network layer reachable information [NLR] messages received from different VPNs.) See [“Configuring a Routing Instance for MPLS VPNs \(CLI Procedure\)” on page 1424](#).

6. For CE routers, configure a routing policy. In addition, if you are not using a route target, configure a VPN routing policy for each PE router in the VPN configuration. Within the policy, describe which packets are sent and received across the VPN and specify how routes are imported into and exported from the router's VRF table. Each advertisement must have an associated route target that uniquely identifies the VPN for which the advertisement is valid. See [“Configuring a Routing Policy for MPLS Layer 3 VPNs \(CLI Procedure\)” on page 1473](#).

RELATED DOCUMENTATION

| [Verifying an MPLS Layer 3 VPN Configuration | 1474](#)

Configuring a Routing Policy for MPLS Layer 3 VPNs (CLI Procedure)

To configure a Layer 3 VPN routing policy on a CE router:

1. Configure the routing policy for the loopback interface.

```
[edit]
user@host# edit policy-options policy-statement policy-name
```

2. Define the term for accepting packets.

```
[edit policy-options policy-statement policy-name]
user@host# set term term-name-accept from protocol direct route-filter local-loopback-address/netmask
exact
user@host# set term term-name-accept then accept
```

3. Define the term for rejecting packets.

```
[edit policy-options policy-statement policy-name]
user@host# set term term-name-reject then reject
```

4. Commit the configuration if you are finished configuring the device.


```
[edit]  
user@host# commit
```

RELATED DOCUMENTATION

| [MPLS Layer 3 VPN Configuration Overview | 1471](#)

Verifying an MPLS Layer 3 VPN Configuration

Purpose

Verify the connectivity of MPLS Layer 3 VPNs using the **ping mpls** command. This command helps to verify that a VPN has been enabled by testing the integrity of the VPN connection between the source and destination routers. The destination prefix corresponds to a prefix in the Layer 3 VPN. However, ping tests only whether the prefix is present in a PE VRF table.

Action

To a combination of a IPv4 destination prefix and a Layer 3 VPN name on the destination PE router, use the following command:

```
ping mpls l3vpn l3vpn-name prefix prefix count count
```

RELATED DOCUMENTATION

| [MPLS Layer 3 VPN Configuration Overview | 1471](#)

Example: Configuring MPLS-Based Layer 3 VPNs

IN THIS SECTION

- [Requirements | 1475](#)
- [Overview and Topology | 1475](#)
- [Configuring the Local PE Switch | 1478](#)
- [Configuring the Remote PE Switch | 1482](#)

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

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 83](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches” on page 88](#).

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 114 on page 1476](#) shows the topology used in this example.

Figure 114: Configuring an MPLS-Based Layer 3 VPN




Table 38 on page 1476 shows the settings of the customer edge interface on the local CE switch.

Table 38: Local CE Switch in the MPLS-Based Layer 3 VPN Topology

Property	Settings	Description
Local CE switch hardware	QFX switch	CE1
Customer edge interface	ge-0/0/14 unit 0 family inet address 51.51.0.14/16 protocols ospf interface ge-0/0/14	Interface that connects CE1 to PE1.

Table 39 on page 1476 shows the settings of the customer edge interface on the remote CE switch.

Table 39: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology

Property	Settings	Description
Remote CE switch hardware	QFX switch	CE2
Customer edge interface	ge-0/0/14 unit 0 family inet address 11.22.26.1/16 protocols ospf interface ge-0/0/14	Interface that connects CE2 to PE2.

Table 40 on page 1476 shows the Layer 3 VPN components of the local PE switch.

Table 40: Layer 3 VPN Components of the Local PE Switch

Property	Settings	Description
Local PE switch hardware	QFX switch	PE1

Table 40: Layer 3 VPN Components of the Local PE Switch (*continued*)

Property	Settings	Description
Customer edge interface	ge-5/0/24 unit 0 family inet address 51.51.0.1/16	Connects PE1 to CE1. NOTE: The family inet 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 family inet address 60.0.0.60/16 family mpls	Connects PE1 to P. NOTE: 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 family inet address 21.21.21.21/32	NOTE: 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 41 on page 1478 shows the Layer 3 VPN components of the remote PE switch.

Table 41: 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 family inet address 11.22.26.14/16	Connects PE2 to CE2. For the Layer 3 VPN configuration, added family mpls . NOTE: The family inet 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 family inet address 60.2.0.60/16 family mpls	Connects PE1 to P. NOTE: 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 family inet address 22.22.22.22/32	NOTE: 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-5/0/24.0
set routing-instances L3VPN-1 protocols ospf interface ge-5/0/24.0
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 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-5/0/24.0
```

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-5/0/24 {
    unit 0 {
      family inet {
        address 51.51.0.1/16;
      }
    }
  }
  xe-0/0/6 {
    unit 0 {
      family inet {
        address 60.0.0.60/16;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 21.21.21.21/32;
      }
    }
  }
  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-options {
        router-id 21.21.21.21;
        autonomous-system 10;
    }
    routing-instances {
        L3VPN-1 {
            instance-type vrf;
            description "BETWEEN PE1 AND PE2";
            route-distinguisher 21:21;
            vrf-target target:21:21;
            vrf-table-label;
            protocols {
                ospf {
                    interface ge-5/0/24.0
                }
            }
        }
    }
}

```

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-0/0/14.0
set routing-instances L3VPN-1 protocols ospf interface ge-0/0/14.0
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-0/0/14.0
```

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-table-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 lo0.0;
            interface xe-0/0/6.0;
        }
    }
}

```



```

    }
  }
  routing-options {
    router-id 22.22.22.22;
    autonomous-system 10;
    routing-instances {
      L3VPN-1 {
        instance-type vrf;
        description "BETWEEN PE1 AND PE2";
        route-distinguisher 21:21;
        vrf-target target:21:21;
        vrf-table-label;
        protocols {
          ospf {
            interface ge-0/0/14.0
          }
        }
      }
    }
  }
}

```

RELATED DOCUMENTATION

[Configuring MPLS on Provider Edge Switches | 83](#)

[Configuring MPLS on Provider Switches | 88](#)

Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

IN THIS SECTION

- [Requirements | 1487](#)
- [Overview | 1487](#)
- [Configuration | 1490](#)
- [Verification | 1497](#)

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

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:

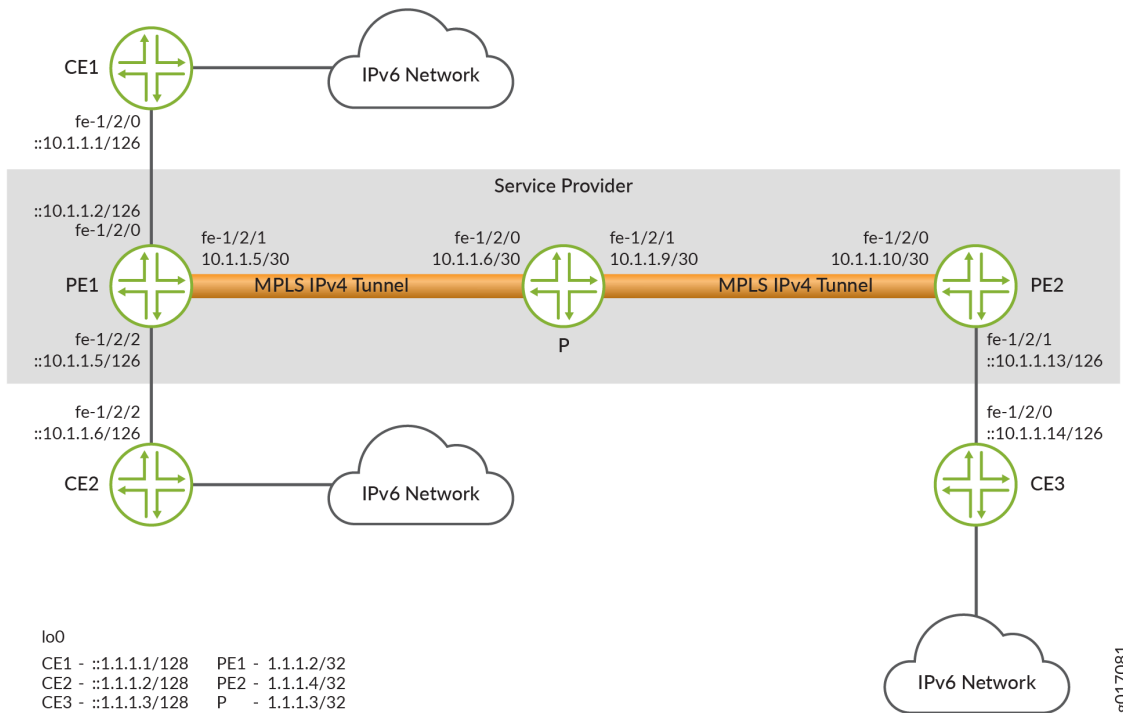
- Internet draft draft-ietf-l3vpn-bgp-ipv6-07.txt, *BGP-MPLS IP VPN extension for IPv6 VPN* (expires January 2006)
- Internet draft draft-ooms-v6ops-bgp-tunnel-06.txt, *Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers* (expires July 2006)

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 17 on page 354](#), 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 115: 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 17 on page 354](#) 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 356](#).

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

Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches

IN THIS SECTION

- [Requirements | 1499](#)
- [Overview and Topology | 1499](#)
- [Configuring the Local PE Switch | 1502](#)
- [Configuring the Remote PE Switch | 1506](#)
- [Verification | 1510](#)

You can implement an MPLS-based Layer 3 virtual private network (VPN) on EX8200 and EX4500 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.

NOTE: The core interfaces and the loopback interfaces are configured in the same way for Layer 2 VPNs and Layer 3 VPNs.

This example shows how to configure an MPLS-based Layer 3 VPN spanning two corporate sites:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for EX Series switches
- Three EX8200 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 Using IP Over MPLS \(CLI Procedure\)” on page 89](#).
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

NOTE: A Layer 3 VPN requires that the PE switches be configured using IP over MPLS.

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.

NOTE: A static path is not configured in this example.

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 116 on page 1500](#) illustrates the topology of this MPLS-based Layer 3 VPN.

Figure 116: MPLS-Based Layer 3 VPN




Table 38 on page 1476 shows the settings of the customer edge interface on the local CE switch.

Table 42: Local CE Switch in the MPLS-Based Layer 3 VPN Topology

Property	Settings	Description
Local CE switch hardware	EX8200 switch	CE1
Customer edge interface	ge-0/0/14 unit 0 family inet address 51.51.0.14/16	Interface that connects CE1 to PE1.

Table 39 on page 1476 shows the settings of the customer edge interface on the remote CE switch.

Table 43: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology

Property	Settings	Description
Remote CE switch hardware	EX8200 switch	CE2
Customer edge interface	ge-0/0/14 unit 0 family inet address 11.22.26.1/16	Interface that connects CE2 to PE2.

Table 40 on page 1476 shows the Layer 3 VPN components of the local PE switch.

Table 44: Layer 3 VPN Components of the Local PE Switch

Property	Settings	Description
Local PE switch hardware	EX8200 switch	PE1

Table 44: Layer 3 VPN Components of the Local PE Switch (*continued*)

Property	Settings	Description
Customer edge interface	ge-5/0/24 unit 0 family inet address 51.51.0.1/16	Connects PE1 to CE1. NOTE: The family inet 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-6/0/0 unit 0 family inet address 60.0.0.60/16 family iso; family mpls	Connects PE1 to P. NOTE: 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 family inet address 21.21.21.21/32 family iso address 49.0001.2102.1021.0210.00	NOTE: 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 41 on page 1478 shows the Layer 3 VPN components of the remote PE switch.

Table 45: Layer 3 VPN Components of the Remote PE Switch

Property	Settings	Description
Remote PE switch hardware	EX8200 switch	PE2
Customer edge interface	ge-11/0/14 unit 0 family inet address 11.22.26.14/16 family mpls	Connects PE2 to CE2. For the Layer 3 VPN configuration, added family mpls . NOTE: The family inet 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-6/0/0/ unit 0 family inet address 60.2.0.60/16 family iso family mpls	Connects PE1 to P. NOTE: 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 family inet address 22.22.22.22/32 family iso address 49.0001.2202.1022.0220.00	NOTE: 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 group ibgp local-address 21.21.21.21 family inet-vpn unicast
set protocols bgp group ibgp type internal
set protocols bgp group ibgp 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-5/0/24.0
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 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 group ibgp 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 ibgp 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 group ibgp 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 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.

7. 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 Configuration Guide](#).

8. 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
```

9. 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> vrf-table-label
```

```
interfaces {
  ge-5/0/24 {
    unit 0 {
      family inet {
        address 51.51.0.1/16;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 21.21.21.21/32;
      }
    }
  }
  xe-6/0/0 {
    unit 0 {
      family inet {
        address 60.0.0.60/16;
      }
      family iso;
      family mpls;
    }
  }
}
protocols {
  mpls {
    label-switched-path 21-22 {
      from 21.21.21.21;
      to 22.22.22.22;
      no-cspf;
    }
    interface xe-6/0/0.0;
    interface lo0.0;
    bgp {
      group ibgp
      type internal
    }
  }
}
```



```

        local-address 21.21.21.21
        family inet-vpn
        unicast
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface ge-5/0/24.0;
            interface lo0.0;
            interface xe-6/0/0.0;
        }
    }
}
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 group ibgp local-address 22.22.22.22 family inet-vpn unicast
set protocols bgp group ibgp type internal
set protocols bgp group ibgp 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 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 group ibgp 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 ibgp 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 group ibgp 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-11/0/14.0
```

7. 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
```

8. 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
```

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@switchPE2# set L3VPN-1 vrf-table-label
```

10. 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-11/0/14 {
    unit 0 {
      family inet {
        address 11.22.26.14/16;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 22.22.22.22/32;
      }
    }
  }
  xe-6/0/0 {
    unit 0 {
      family inet {
        address 60.2.0.60/16;
      }
      family iso;
      family mpls;
    }
  }
}
protocols {
```



```

mpls {
    label-switched-path 22-21 {
        from 22.22.22.22;
        to 21.21.21.21;
        no-cspf;
    }
    interface xe-6/0/0.0;
    interface lo0.0;
    bgp {
        group ibgp
        type internal
        local-address 21.21.21.21
        family inet-vpn
        unicast
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface ge-11/0/14.0;
            interface lo0.0;
            interface xe-6/0/0.0;
        }
    }
}
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;
}

```


Verification

IN THIS SECTION

- [Verifying Peering and Adjacency | 1510](#)
- [Verifying That the Local CE Switch Can Ping the Local PE Switch | 1510](#)
- [Verifying That the Local PE Switch Can Ping the Local CE Switch | 1511](#)

To confirm that the MPLS-based Layer 3 VPN is working properly, perform these tasks:

Verifying Peering and Adjacency

Purpose

Verify the peering and adjacency along the route from CE1 (the local CE switch or router) to CE2 (the remote CE switch or router), starting with checking the routing protocol adjacency on the local PE switch:

NOTE: Be sure to specify the name of the routing instance.

Action

```
user@switchPE1> show ospf neighbor instance L3VPN-1
```

Address	Interface	State	ID	Pri	Dead
51.51.0.14	ge-5/0/24.0	Full	21.21.21.21	128	38

Meaning

The **Address** field shows the IP address of the customer edge interface that connects CE1 to PE1. The **Interface** field shows the interface name of the customer edge interface that connects PE1 to CE1. For our purposes, the **State** field is the most important. It shows a status of **Full**, indicating that neighboring routing devices are fully adjacent. These adjacencies appear in router-link and network-link advertisements. (The field **Pri** indicates the priority of the neighbor to become the designated router. The field **Dead** indicates the number of seconds until the neighbor becomes unreachable.)

Verifying That the Local CE Switch Can Ping the Local PE Switch

Purpose

Verify that the local CE switch can ping the local PE switch:

Action

```
user@switchCE1> ping 51.51.0.1
```

```
PING 51.51.0.1 (51.51.0.1): 56 data bytes
64 bytes from 51.51.0.1: icmp_seq=0 ttl=64 time=3.461 ms
64 bytes from 51.51.0.1: icmp_seq=1 ttl=64 time=3.543 ms
```

Meaning

This command specified the IP address of the customer edge interface on PE1. The results indicate that CE1 is receiving packets from PE1.

Verifying That the Local PE Switch Can Ping the Local CE Switch

Purpose

Verify that the local PE switch can ping the local CE switch:

Action

```
user@switchPE1> ping 51.51.0.14 routing-instance L3VPN-1
```

```
PING 51.51.0.14 (51.51.0.14): 56 data bytes
64 bytes from 51.51.0.14: icmp_seq=0 ttl=64 time=3.842 ms
64 bytes from 51.51.0.14: icmp_seq=1 ttl=64 time=3.736 ms
```

Meaning

The results indicate a successful connection.

RELATED DOCUMENTATION

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

Configuring an MPLS-Based Layer 3 VPN (CLI Procedure)

You can configure MPLS-based Layer 3 virtual private networks (VPNs) on EX8200 and EX4500 switches. Layer 3 VPNs leverage the service provider's technical expertise for site-to-site routing.

To configure Layer 3 VPN functionality in your MPLS network, you must enable Layer 3 VPN support on the local and remote provider edge (PE) switches as described in this task.

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 Using IP Over MPLS \(CLI Procedure\)” on page 89](#).
- Configure one or more provider switches. See [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

NOTE: A Layer 3 VPN requires that the PE switches be configured using IP over MPLS.

Configure the Layer 3 VPN components on both PE switches. This procedure describes how to configure one PE switch. Repeat the procedure to configure the remote PE switch.

NOTE: When you configure the remote PE switch, the information specified for the routing instance must be configured the same as the information specified for the routing instance on the local PE switch. You must also specify the same BGP group name. The following statements will have different values on the remote PE switch from those on the local PE switch:

- **local-address**
- **neighbor**

To configure an MPLS-based Layer 3 VPN on the PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:

```
[edit protocols bgp]
user@switch# set local-address address family inet-vpn unicast
```

2. Configure the BGP group, specifying the group name and **type internal**:

```
[edit protocols bgp]
user@switch# set group group-name 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@switch# set neighbor address
```

4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:

```
[edit]
user@switch# set routing-instances routing-instance-name instance-type vrf
```

5. Configure a description for this routing instance:

```
[edit]
user@switch# set routing-instances routing-instance-name description text
```

6. Configure the routing instance to use a route distinguisher:

NOTE: Each routing instance that you configure on a PE switch must have a unique route distinguisher associated with it. VPN routing instances must have 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.

```
user@switch# set routing-instances routing-instance-name route-distinguisher ip-address:number
```

7. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@switch# set routing-instance-name vrf-target community
```

NOTE: If you configure the **community** option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. 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 Configuration Guide](#).

8. 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@switch# set routing-instance-name vrf-table-label
```

9. (Optional) Configure the routing options:

NOTE: We recommend that you 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@switch# set router-id ip-address autonomous-system as-number
```

RELATED DOCUMENTATION

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

[Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches | 1456](#)

Configuring TCC and Layer 2.5 Switching

IN THIS CHAPTER


- TCC and Layer 2.5 Switching Overview | 1515
- Configuring VLAN TCC Encapsulation | 1516
- Configuring Translation Cross-Connect Interface Switching | 1518

TCC and Layer 2.5 Switching Overview

Translational cross-connect (TCC) allows you to forward traffic between a variety of Layer 2 protocols or circuits. It is similar to its predecessor, CCC. However, while CCC requires the same Layer 2 encapsulations on both sides of a router (such as Point-to-Point Protocol [PPP] or Frame Relay-to-Frame Relay), TCC lets you connect different types of Layer 2 protocols interchangeably. With TCC, combinations such as PPP-to-ATM and Ethernet-to-Frame Relay cross-connections are possible. Also, TCC can be used to create Layer 2.5 VPNs and Layer 2.5 circuits.

Consider a sample topology ([Figure 117 on page 1515](#)) in which you can configure a full-duplex Layer 2.5 translational cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the TCC interface. In this topology, Router B strips all PPP encapsulation data from frames arriving from Router A and adds ATM encapsulation data before the frames are sent to Router C. All Layer 2 negotiations are terminated at the interconnecting router (Router B).

Figure 117: Sample Translation Cross-Connect Topology



TCC functionality is different from standard Layer 2 switching. TCC only swaps Layer 2 headers. No other processing, such as header checksums, time-to-live (TTL) decrementing, or protocol handling, is performed. Currently, TCC is supported in IPv4, ISO, and MPLS.

Ethernet TCC is supported on interfaces that carry IPv4 traffic only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC and extended VLAN CCC are not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC and extended VLAN TCC are not supported.

RELATED DOCUMENTATION

[Configuring VLAN TCC Encapsulation | 1516](#)

[Configuring Translation Cross-Connect Interface Switching | 1518](#)

Ethernet Interfaces User Guide for Routing Devices

Configuring VLAN TCC Encapsulation

VLAN TCC encapsulation allows circuits to have different media on either side of the forwarding path. VLAN TCC encapsulation supports TPID 0x8100 only. You must include configuration statements at the logical and physical interface hierarchy levels.

To configure VLAN TCC encapsulation, include the **encapsulation** statement and specify the **vlan-tcc** option:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-tcc;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]**

Additionally, configure the logical interface by including the **proxy** and **remote** statements:

```
proxy {
  inet-address;
}
remote {
  (inet-address | mac-address);
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family tcc]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family tcc]**

The proxy address is the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy.

The remote address is the IP or MAC address of the remote router. The **remote** statement provides ARP capability from the TCC switching router to the Ethernet neighbor. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

When VLAN TCC encapsulation is configured on the logical interface, you also must specify flexible Ethernet services on the physical interface. To specify flexible Ethernet services, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level and specify the **flexible-ethernet-services** option:

```
[edit interfaces interface-name]  
encapsulation flexible-ethernet-services;
```

Extended VLAN TCC encapsulation supports TPIDs 0x8100 and 0x9901. Extended VLAN TCC is specified at the physical interface level. When configured, all units on that interface must use VLAN TCC encapsulation, and no explicit configuration is needed on logical interfaces.

One-port Gigabit Ethernet, 2-port Gigabit Ethernet, and 4-port Fast Ethernet PICs with VLAN tagging enabled can use VLAN TCC encapsulation. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the **[edit interfaces *interface-name*]** hierarchy level and specify the **extended-vlan-tcc** option:

```
[edit interfaces interface-name]  
encapsulation extended-vlan-tcc;
```

For VLAN TCC encapsulation, all VLAN IDs from 1 through 1024 are valid. VLAN ID 0 is reserved for tagging the priority of frames.

Extended VLAN TCC is not supported on 4-port Gigabit Ethernet PICs.

RELATED DOCUMENTATION

[encapsulation](#) | [3007](#)

[remote](#)

[proxy](#)

[TCC and Layer 2.5 Switching Overview](#) | [1515](#)

[Configuring Translation Cross-Connect Interface Switching](#) | [1518](#)

[Ethernet Interfaces User Guide for Routing Devices](#)

Configuring Translation Cross-Connect Interface Switching

To configure a full-duplex Layer 2.5 translation cross-connect between two routers (A and C), you can configure a Juniper Networks router (Router B) as the TCC interface. Ethernet TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. Consider the topology in [Figure 118 on page 1518](#) where the Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit accepts packets carrying standard TPID values.

Figure 118: Sample Topology of Layer 2.5 Translational Cross-Connect



If traffic flows from Router A to Router C, the Junos OS strips all PPP encapsulation data from incoming packets and adds Ethernet encapsulation data before forwarding the packets. If traffic flows from Router C to Router A, the Junos OS strips all Ethernet encapsulation data from incoming packets and adds PPP encapsulation data before forwarding the packets.

To configure the router as the translational cross-connect interface:

1. In the configuration mode, at the **[edit]** hierarchy level, first configure the interface that is connected to Router A.

```
[edit]
user@host# edit interfaces interface-name
```

2. (Optional) Specify the description of the interface. For example, you could specify the interface name on Router A that is connected to this interface.

```
[edit interfaces interface-name]
user@host# set description description
```

3. Specify the encapsulation. If the Router A to Router B circuit is PPP, then specify **ppp-tcc** as the encapsulation. If the Router A to Router B circuit is frame relay, specify **frame-relay-tcc**.

```
[edit interfaces interface-name]
user@host# set encapsulation encapsulation-type
```

4. In the configuration mode, at the **[edit]** hierarchy level, first configure the interface that is connected to Router C.


```
[edit]
user@host# edit interfaces interface-name
```

5. (Optional) Specify the description of this interface. For example, you could specify the interface name on Router C that is connected to this interface.

```
[edit interfaces interface-name]
user@host# set description description
```

6. Specify the encapsulation. If the Router B to Router C circuit is Ethernet, then specify **ethernet-tcc** as the encapsulation. If the Router B to Router C circuit is ATM, specify **atm-tcc-vc-mux**.

```
[edit interfaces interface-name]
user@host# set encapsulation encapsulation-type
```

7. Specify the IP address or MAC address of the remote router to provide address resolution protocol (ARP) for the TCC router's Ethernet-based neighbor using the **remote** statement. You must specify the statement at the **[edit interfaces interface-name unit unit-number family tcc]** hierarchy level. You can also specify the MAC address of the remote router instead of the IP address. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

```
[edit interfaces interface-name]
user@host# set unit 0 family family remote inet-address ip-address
```

8. Specify the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy using the **proxy** statement. You must specify the statement at the **[edit interfaces interface-name unit unit-number family tcc]** hierarchy level.

```
[edit interfaces interface-name]
user@host# set unit 0 family family proxy inet-address ip-address
```

To verify the TCC connection, use the **show connections** command on TCC router.

RELATED DOCUMENTATION

[encapsulation](#) | [3007](#)

[remote](#)

proxy

[TCC and Layer 2.5 Switching Overview](#) | 1515

[Configuring VLAN TCC Encapsulation](#) | 1516

Ethernet Interfaces User Guide for Routing Devices

Configuring CLNS VPNs

IN THIS CHAPTER

- [CLNS Overview | 1521](#)
- [CLNS Configuration Overview | 1522](#)
- [Understanding ES-IS for CLNS | 1524](#)
- [Example: Configuring ES-IS for CLNS | 1525](#)
- [Understanding IS-IS for CLNS | 1527](#)
- [Example: Configuring IS-IS for CLNS | 1528](#)
- [Understanding Static Routes for CLNS | 1531](#)
- [Example: Configuring Static Routes for CLNS When No IGP is Present | 1532](#)
- [Understanding BGP for CLNS VPNs | 1535](#)
- [Example: Configuring BGP for CLNS VPNs | 1535](#)
- [Example: Configuring a VPN Routing Instance for CLNS | 1538](#)
- [Verifying a CLNS VPN Configuration | 1541](#)

CLNS Overview

Connectionless Network Service (CLNS) is a Layer 3 protocol similar to IP version 4 (IPv4) for linking hosts (end systems) with routers (intermediate systems) in an Open Systems Interconnection (OSI) network. CLNS and its related OSI protocols, Intermediate System-to-Intermediate System (IS-IS) and End System-to-Intermediate System (ES-IS), are International Organization for Standardization (ISO) standards.

You can configure devices running Junos OS as provider edge (PE) routers within a CLNS network. CLNS networks can be connected over an IP MPLS network core using Border Gateway Protocol (BGP) and MPLS Layer 3 virtual private networks (VPNs). See RFC 2547, *BGP/MPLS VPNs*.

CLNS uses network service access points (NSAPs), similar to IP addresses found in IPv4, to identify end systems (hosts) and intermediate systems (routers). ES-IS enables the hosts and routers to discover each other. IS-IS is the interior gateway protocol (IGP) that carries ISO CLNS routes through a network.

For more information about CLNS, see the ISO 8473 standards.

RELATED DOCUMENTATION

- | |
|---|
| CLNS Configuration Overview 1522 |
| Understanding ES-IS for CLNS 1524 |
| Understanding IS-IS for CLNS 1527 |
| Understanding Static Routes for CLNS 1531 |
| Understanding BGP for CLNS VPNs 1535 |

CLNS Configuration Overview

To configure CLNS:

1. Configure the network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.
2. If applicable, configure BGP and VPNs. See:
 - [Example: Configuring BGP for CLNS VPNs on page 1535](#)
 - [MPLS Layer 2 VPN Configuration Overview on page 1429](#)
 - [MPLS Layer 3 VPN Configuration Overview on page 1471](#)

3. Configure a VPN routing instance. You typically configure ES-IS, IS-IS, and CLNS static routes using a VPN routing instance. See [“Example: Configuring a VPN Routing Instance for CLNS” on page 1538](#).
4. Configure one or more of the following protocols for CLNS (depending on your network).
 - ES-IS—If a device is a PE router within a CLNS island that contains any end systems, you must configure ES-IS on the device. If a CLNS island does not contain any end systems, you do not need to configure ES-IS on a device. See [“Example: Configuring ES-IS for CLNS” on page 1525](#).

NOTE: ES-IS is enabled only if either ES-IS or IS-IS is configured on the router. ES-IS must not be disabled. If ES-IS is not explicitly configured, the interface sends and receives only intermediate system hello (ISH) messages. If ES-IS is explicitly configured and disabled, the interface does not send or receive ES-IS packets. If ES-IS is explicitly configured and not disabled, the interface sends and receives ISH messages as well as ES-IS packets.

One of the interfaces that is configured for ES-IS must be configured with an ISO address for hello messages. The ISO address family must be configured on an interface to support ES-IS on that interface.

- IS-IS—You can configure IS-IS to exchange CLNS routes within a CLNS island. See [“Example: Configuring IS-IS for CLNS” on page 1528](#).

NOTE: If you have a pure CLNS island—an island that does not contain any IP devices—you must disable IPv4 and IPv6 routing. Also, to export BGP routes into IS-IS, you must configure and apply an export policy.

- Static routes—If some devices in your network do not support IS-IS, you must configure CLNS static routes. You can use static routing with or without IS-IS. You might also consider using static routes if your network is simple. See [“Example: Configuring Static Routes for CLNS When No IGP is Present” on page 1532](#).
- BGP—See [“Example: Configuring BGP for CLNS VPNs” on page 1535](#).

NOTE: Many of the configuration statements used to configure CLNS and routing protocols can be included at different hierarchy levels in the configuration.

RELATED DOCUMENTATION

[CLNS Overview | 1521](#)

[Verifying a CLNS VPN Configuration | 1541](#)

Understanding ES-IS for CLNS

End System-to-Intermediate System (ES-IS) is a protocol that resolves Layer 3 ISO network service access points (NSAP) to Layer 2 addresses. ES-IS has an equivalent role as Address Resolution Protocol (ARP) in IP version 4 (IPv4).

ES-IS provides the basic interaction between Connectionless Network Service (CLNS) hosts (end systems) and routers (intermediate systems). ES-IS allows hosts to advertise NSAP addresses to other routers and hosts attached to the network. Those routers can then advertise the address to the rest of the network by using Intermediate System-to-Intermediate System (IS-IS). Routers use ES-IS to advertise their network entity title (NET) to hosts and routers that are attached to that network.

ES-IS routes are exported to Layer 1 IS-IS by default. You can also export ES-IS routes into Layer 2 IS-IS by configuring a routing policy. ES-IS generates and receives end system hello (ESH) hello messages when the protocol is configured on an interface. ES-IS is a resolution protocol that allows a network to be fully ISO integrated at both the network layer and the data layer.

The resolution of Layer 3 ISO NSAPs to Layer 2 subnetwork point of attachments (SNPAs) by ES-IS is equivalent to ARP within an IPv4 network. If a device is a provider edge (PE) router within a CLNS island that contains any end systems, you must configure ES-IS on the device.

For more information about ES-IS, see the ISO 9542 standard.

RELATED DOCUMENTATION

[CLNS Overview | 1521](#)

[Example: Configuring ES-IS for CLNS | 1525](#)

Example: Configuring ES-IS for CLNS

IN THIS SECTION

- [Requirements | 1525](#)
- [Overview | 1525](#)
- [Configuration | 1525](#)
- [Verification | 1527](#)

This example shows how to create a routing instance and enable ES-IS for CLNS on all interfaces.

Requirements

Before you begin, configure the network interfaces. See *Interfaces User Guide for Security Devices*.

Overview

The configuration instructions in this topic describe how to create a routing-instance called `aaaa`, set the end system configuration timer for the interfaces to 180, and set a preference value to 30 for ES-IS.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances aaaa protocols esis interface all end-system-configuration-timer 180
set routing-instances aaaa protocols esis preference 30
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure ES-IS for CLNS:

1. Configure the routing instance.


```
[edit]
user@host# edit routing-instances aaaa
```

2. Enable ES-IS on all interfaces.

```
[edit routing-instances aaaa]
user@host# set protocols esis interface all
```

3. Configure the end system configuration timer.

```
[edit routing-instances aaaa]
user@host# set protocols esis interface all end-system-configuration-timer 180
```

4. Configure the preference value.

```
[edit routing-instances aaaa]
user@host# set protocols esis preference 30
```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show routing-instances
aaaa {
  protocols {
    esis {
      preference 30;
      interface all {
        end-system-configuration-timer 180;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying Routing-Instance for CLNS | 1527](#)
- [Verifying ES-IS for CLNS | 1527](#)

Confirm that the configuration is working properly.

Verifying Routing-Instance for CLNS

Purpose

Verify that the policy options are enabled for the routing instance.

Action

From operational mode, enter the **show routing-instances** command.

Verifying ES-IS for CLNS

Purpose

Verify that ES-IS is enabled.

Action

From operational mode, enter the **show protocols** command.

RELATED DOCUMENTATION

[CLNS Configuration Overview | 1522](#)

[Understanding ES-IS for CLNS | 1524](#)

[Verifying a CLNS VPN Configuration | 1541](#)

Understanding IS-IS for CLNS

IS-IS extensions provide the basic interior gateway protocol (IGP) support for collecting intradomain routing information for Connectionless Network Service (CLNS) destinations within a CLNS network. Routers that learn host addresses through End System-to-Intermediate System (ES-IS) can advertise the addresses to other routers (intermediate systems) by using IS-IS.

For more information about IS-IS, see the ISO 10589 standard.

RELATED DOCUMENTATION

[CLNS Overview | 1521](#)

[Example: Configuring IS-IS for CLNS | 1528](#)

Example: Configuring IS-IS for CLNS

IN THIS SECTION

- [Requirements | 1528](#)
- [Overview | 1528](#)
- [Configuration | 1528](#)
- [Verification | 1530](#)

This example shows how to create a routing instance and enable IS-IS protocol on all interfaces.

Requirements

Before you begin, configure the network interfaces. See [Interfaces User Guide for Security Devices](#).

Overview

The configuration instructions in this topic describe how to create a routing-instance called `aaaa`, enable IS-IS on all interfaces, and define BGP export policy name (`dist-bgp`), family (ISO), and protocol (BP), and apply the export policy to IS-IS.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter **commit** from configuration mode.


```

set routing-instances aaaa protocols isis clns-routing
set routing-instances aaaa protocols isis interface all
set routing-instances aaaa protocols isis no-ipv4-routing no-ipv6-routing
set policy-options policy-statement dist-bgp from family iso protocol bgp
set policy-options policy-statement dist-bgp then accept
set routing-instances aaaa protocols isis export dist-bgp

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure IS-IS for CLNS:

1. Configure the routing instance.

```

[edit]
user@host# edit routing-instances aaaa

```

2. Enable CLNS routing.

```

[edit routing-instances aaaa]
user@host# set protocols isis clns-routing

```

3. Enable IS-IS on all interfaces.

```

[edit routing-instances aaaa]
user@host# set protocols isis interface all

```

4. (Optional) Disable IPv4 and IPv6 routing to configure a pure CLNS network .

```

[edit routing-instances aaaa]
user@host# set protocols isis no-ipv4-routing no-ipv6-routing

```

5. Define the BGP export policy name, family, and protocol.

```

[edit policy-options]
user@host# set policy-statement dist-bgp from family iso protocol bgp

```

6. Define the action for the export policy.


```
[edit policy-options]
user@host# set policy-statement dist-bgp then accept
```

7. Apply the export policy to IS-IS.

```
[edit routing-instances aaaa]
user@host# set protocols isis export dist-bgp
```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show routing-instances
aaaa {
  protocols {
    isis {
      export dist-bgp;
      no-ipv4-routing;
      no-ipv6-routing;
      clns-routing;
      interface all;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying Routing-Instance for CLNS | 1531](#)
- [Verifying IS-IS for CLNS | 1531](#)

Confirm that the configuration is working properly.

Verifying Routing-Instance for CLNS

Purpose

Verify that the policy options are enabled for the routing instance.

Action

From operational mode, enter the **show routing-instances** command.

Verifying IS-IS for CLNS

Purpose

Verify that IS-IS is enabled.

Action

From operational mode, enter the **show protocols** command.

RELATED DOCUMENTATION

[CLNS Configuration Overview | 1522](#)

[Understanding IS-IS for CLNS | 1527](#)

[Verifying a CLNS VPN Configuration | 1541](#)

Understanding Static Routes for CLNS

The Connectionless Network Service (CLNS) is an ISO Layer 3 protocol that uses network service access point (NSAP) reachability information instead of IPv4 or IPv6 prefixes.

You can configure static routes to exchange CLNS routes within a CLNS island. A *CLNS island* is typically an IS-IS level 1 area that is part of a single IGP routing domain. An island can contain more than one area. CLNS islands can be connected by VPNs.

RELATED DOCUMENTATION

[Example: Configuring Static Routes for CLNS When No IGP is Present | 1532](#)

Example: Configuring Static Routes for CLNS When No IGP is Present

IN THIS SECTION

- [Requirements | 1532](#)
- [Overview | 1532](#)
- [Configuration | 1532](#)
- [Verification | 1534](#)

This example shows how to configure static routes for CLNS.

Requirements

Before you begin, configure the network interfaces. See *Interfaces User Guide for Security Devices*.

Overview

In this example, you configure static routes for CLNS. In the absence of an interior gateway protocol (IGP) on a certain link, a routing device might need to be configured with static routes for CLNS prefixes to be reachable by way of that link. This might be useful, for example, at an autonomous system (AS) boundary.

When you configure static routes for CLNS, consider the following tasks:

- Specify the **iso.0** routing table option to configure a primary instance CLNS static route.
- Specify the **instance-name.iso.0** routing table option to configure a CLNS static route for a particular routing instance.
- Specify the **route nsap-prefix** statement to configure the destination for the CLNS static route.
- Specify the **next-hop (interface-name | iso-net)** statement to configure the next hop, specified as an ISO network entity title (NET) or interface name.
- Include the **qualified-next-hop (interface-name | iso-net)** statement to configure a secondary backup next hop, specified as an ISO network entity title or interface name.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4212
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152 next-hop
  t1-0/2/2.0
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.0000.0000.0000/152 qualified-next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 preference 20
set routing-options rib iso.0 static iso-route 47.0005.80ff.f800.0000.0000.0000.0000/152 qualified-next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 metric 10
```

Step-by-Step Procedure

To configure static routes for CLNS:

1. Configure the routes.

```
[edit routing-options rib iso.0 static]
user@host# set iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4212
user@host# set iso-route 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152 next-hop t1-0/2/2.0
user@host# set iso-route 47.0005.80ff.f800.0000.0000.0000.0000/152 qualified-next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 preference 20
user@host# set iso-route 47.0005.80ff.f800.0000.0000.0000.0000/152 qualified-next-hop
  47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 metric 10
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Results

Confirm your configuration by issuing the **show routing-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options
rib iso.0 {
  static {
    iso-route 47.0005.80ff.f800.0000.ffff.ffff/152 next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212;
    iso-route 47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152 next-hop t1-0/2/2.0;
```



```

iso-route 47.0005.80ff.f800.0000.eee0/152 {
    qualified-next-hop 47.0005.80ff.f800.0000.0108.0001.1921.6800.4002 {
        preference 20;
        metric 10;
    }
}
}
}

```

Verification

Checking the Routing Table

Purpose

Make sure that the expected routes appear in the routing table.

Action

```
user@host> show route table iso.0
```

```

iso.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1921.6800.4212/152
    *[Static/5] 00:00:25
    > via t1-0/2/2.0
47.0005.80ff.f800.0000.eee0/84
    *[Static/20] 00:04:01, metric 10, metric2 10
    > to #75 0.12.0.34.0.56 via fe-0/0/1.0
47.0005.80ff.f800.0000.ffff.ffff/104
    *[Static/5] 00:04:01, metric2 0
    > via t1-0/2/2.0

```

Meaning

The static routes appear in the routing table.

RELATED DOCUMENTATION

[CLNS Configuration Overview | 1522](#)

[Understanding Static Routes for CLNS | 1531](#)

Understanding BGP for CLNS VPNs

BGP extensions allow BGP to carry Connectionless Network Service (CLNS) virtual private network (VPN) network layer reachability information (NLRI) between provider edge (PE) routers. Each CLNS route is encapsulated into a CLNS VPN NLRI and propagated between remote sites in a VPN.

CLNS is a Layer 3 protocol similar to IP version 4 (IPv4). CLNS uses network service access points (NSAPs) to address end systems. This allows for a seamless autonomous system (AS) based on International Organization for Standardization (ISO) NSAPs.

A single routing domain consisting of ISO NSAP devices are considered to be CLNS islands. CLNS islands are connected together by VPNs.

You can configure BGP to exchange ISO CLNS routes between PE routers connecting various CLNS islands in a VPN using multiprotocol BGP extensions. These extensions are the ISO VPN NLRIs.

Each CLNS network island is treated as a separate VPN routing and forwarding instance (VRF) instance on the PE router.

You can configure CLNS on the global level, group level, and neighbor level.

RELATED DOCUMENTATION

[CLNS Overview | 1521](#)

[Example: Configuring BGP for CLNS VPNs | 1535](#)

Example: Configuring BGP for CLNS VPNs

IN THIS SECTION

- [Requirements | 1536](#)
- [Overview | 1536](#)
- [Configuration | 1536](#)
- [Verification | 1537](#)

This example shows how to create a BGP group for CLNS VPNs, define the BGP peer neighbor address for the group, and define the family.

Requirements

Before you begin, configure the network interfaces. See the *Interfaces User Guide for Security Devices*.

Overview

In this example, you create the BGP group called pedge-pedge, define the BGP peer neighbor address for the group as 10.255.245.215, and define the BGP family.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set protocols bgp group pedge-pedge neighbor 10.255.245.213
set protocols bgp family iso-vpn unicast
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure BGP for CLNS VPNs:

1. Configure the BGP group and define the BGP peer neighbor address.

```
[edit protocols bgp]
user@host# set group pedge-pedge neighbor 10.255.245.213
```

2. Define the family.

```
[edit protocols bgp]
user@host# set family iso-vpn unicast
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```


Verification

Verifying the Neighbor Status

Purpose

Display information about the BGP peer.

Action

From operational mode, run the **show bgp neighbor 10.255.245.213** command. Look for **iso-vpn-unicast** in the output.

```
user@host> show bgp neighbor 10.255.245.213
```

```
Peer: 10.255.245.213+179 AS 200 Local: 10.255.245.214+3770 AS 100
Type: External State: Established Flags: <ImportEval Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: None
Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
Rib-group Refresh>
Address families configured: iso-vpn-unicast
Local Address: 10.255.245.214 Holdtime: 90 Preference: 170
Number of flaps: 0
Peer ID: 10.255.245.213 Local ID: 10.255.245.214 Active Holdtime: 90
Keepalive Interval: 30 Peer index: 0
NLRI advertised by peer: iso-vpn-unicast
NLRI for this session: iso-vpn-unicast
Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: in sync
Active prefixes: 3
Received prefixes: 3
Suppressed due to damping: 0
Advertised prefixes: 3
Table aaaa.iso.0
RIB State: BGP restart is complete
RIB State: VPN restart is complete
Send state: not advertising
Active prefixes: 3
Received prefixes: 3
Suppressed due to damping: 0
Last traffic (seconds): Received 6 Sent 5 Checked 5
Input messages: Total 1736 Updates 4 Refreshes 0 Octets 33385
```



```
Output messages: Total 1738 Updates 3 Refreshes 0 Octets 33305
Output Queue[0]: 0
Output Queue[1]: 0
```

RELATED DOCUMENTATION

[CLNS Configuration Overview | 1522](#)

[Understanding BGP for CLNS VPNs | 1535](#)

[Verifying a CLNS VPN Configuration | 1541](#)

Example: Configuring a VPN Routing Instance for CLNS

IN THIS SECTION

- [Requirements | 1538](#)
- [Overview | 1538](#)
- [Configuration | 1539](#)
- [Verification | 1540](#)

This example shows how to create a CLNS routing instance and set the instance type for Layer 3 VPNs.

Requirements

Before you begin, configure the network interfaces. See *Interfaces User Guide for Security Devices*.

Overview

The following example shows how to create a CLNS routing instance called `aaaa` and set the instance type to VRF for Layer 3 VPNs. Within the example, you specify that the `lo0.1` interface, `e1-2/0/0.0` interface, and `t1-3/0/0.0` interface all belong to the routing instance. The route distinguisher is set as `10.255.245.1:1` and the policy for the Layer 3 VRF table is set as `target:11111:1`.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances aaaa instance-type vrf
set routing-instances aaaa interface lo0.1
set routing-instances aaaa interface ge-0/0/3
set routing-instances aaaa interface ge-0/0/2
set routing-instances aaaa route-distinguisher 10.255.245.1:1
set routing-instances aaaa vrf-target target:1111:1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a VPN routing instance:

1. Create the routing instance.

```
[edit]
user@host# edit routing-instances aaaa
```

2. Specify the routing instance type.

```
[edit routing-instances aaaa]
user@host# set instance-type vrf
```

3. Specify the interfaces that belong to the routing instance.

```
[edit routing-instances aaaa]
user@host# set interface lo0.1
user@host# set interface ge-0/0/3
user@host# set interface ge-0/0/2
```

4. Specify the route distinguisher.

```
[edit routing-instances aaaa]
user@host# set route-distinguisher 10.255.245.1:1
```


5. Specify the policy for the Layer 3 VRF table.

```
[edit routing-instances aaaa]
user@host# set vrf-target target:11111:1
```

6. Enable family ISO on the interfaces edit interfaces interface-name unit-id.

```
[edit routing-instances aaaa]
user@host# set family ISO
```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit ]
user@host# show routing-instances
instance-type vrf;
interface ge-0/0/2.0;
interface ge-0/0/3.0;
interface lo0.1;
route-distinguisher 10.255.245.1:1;
vrf-target target:11111:1;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Configured CLNS Routing Instance

Purpose

Confirm that the configuration is working properly.

Verify that the CLNS routing instance is configured.

Action

From operational mode, enter the **show routing-instances** command.

RELATED DOCUMENTATION

[CLNS Configuration Overview](#) | 1522

Verifying a CLNS VPN Configuration

Purpose

Verify that the device is configured correctly for CLNS VPNs.

Action

From configuration mode in the CLI, enter the **show** command.

```
[edit]
user@host# show
interfaces {
  e1-2/0/0.0 {
    unit 0 {
      family inet {
        address 192.168.37.51/31;
      }
      family iso;
      family mpls;
    }
  }
  t1-3/0/0.0 {
    unit 0 {
      family inet {
        address 192.168.37.24/32;
      }
      family iso;
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 127.0.0.1/32;
        address 10.255.245.215/32;
      }
      family iso {
        address 47.0005.80ff.f800.0000.0108.0001.1921.6800.4215.00;
      }
    }
    unit 1 {
```



```

        family iso {
            address 47.0005.80ff.f800.0000.0108.aaa2.1921.6800.4215.00;
        }
    }
}
routing-options {
    autonomous-system 230;
}
protocols {
    bgp {
        group pedge-pegde {
            type internal;
            local-address 10.255.245.215;
            neighbor 10.255.245.212 {
                family iso-vpn {
                    unicast;
                }
            }
        }
    }
}
policy-options {
    policy-statement dist-bgp {
        from {
            protocol bgp;
            family iso;
        }
        then accept;
    }
}
routing-instances {
    aaaa {
        instance-type vrf;
        interface lo0.1;
        interface e1-2/0/0.0;
        interface t1-3/0/0.0;
        route-distinguisher 10.255.245.1:1;
        vrf-target target:11111:1;
        routing-options {
            rib aaaa.iso.0 {
                static {
                    iso-route 47.0005.80ff.f800.0000.bbbb.1022/104
                        next-hop 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00;
                }
            }
        }
    }
}

```



```
    }  
  }  
}  
protocols {  
  esis {  
    interface all;  
  }  
  isis {  
    export dist-bgp;  
    no-ipv4-routing;  
    no-ip64-routing;  
    clns-routing;  
    interface all;  
  }  
}  
}
```

RELATED DOCUMENTATION

[CLNS Configuration Overview](#) | 1522

Configuring VPLS

IN THIS CHAPTER

- VPLS Overview | 1546
- VPLS Configuration Overview | 1550
- Migrating from FEC128 LDP-VPLS to EVPN Overview | 1552
- Understanding VPLS Interfaces | 1561
- Example: Configuring Routing Interfaces on the VPLS PE Router | 1564
- Example: Configuring the Interface to the VPLS CE Device | 1565
- VPLS Filters and Policers Overview | 1567
- Example: Configuring VPLS Filters | 1567
- Example: Configuring VPLS Policers | 1571
- Understanding VPLS Routing Instances | 1573
- Example: Configuring the VPLS Routing Instance | 1576
- Example: Configuring Automatic Site Identifiers for VPLS | 1580
- Example: Configuring OSPF on the VPLS PE Router | 1581
- Example: Configuring RSVP on the VPLS PE Router | 1583
- Example: Configuring MPLS on the VPLS PE Router | 1585
- Example: Configuring LDP on the VPLS PE Router | 1587
- Example: Configuring VPLS over GRE with IPsec VPNs | 1588
- Example: Configuring VPLS with BGP Signaling | 1610
- Example: Configuring BGP on the VPLS PE Router | 1628
- Example: Configuring Routing Options on the VPLS PE Router | 1630
- Understanding VPLS VLAN Encapsulation | 1631
- Understanding VPLS VLAN Encapsulation on a Logical Interface | 1632
- Example: Configuring VPLS VLAN Encapsulation | 1633
- Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces | 1637
- Example: Configuring Extended VLAN VPLS Encapsulation | 1639

VPLS Overview

IN THIS SECTION

- [Sample VPLS Topology | 1546](#)
- [VPLS on PE Routers | 1547](#)
- [Using an Ethernet Switch as the VPLS CE Device | 1549](#)
- [VPLS Exceptions on SRX Series Devices | 1549](#)

Virtual private LAN service (VPLS) is an Ethernet-based point-to-multipoint Layer 2 VPN. It allows you to connect geographically dispersed Ethernet LAN sites to each other across an MPLS backbone. For customers who implement VPLS, all sites appear to be in the same Ethernet LAN even though traffic travels across the service provider's network.

VPLS, in its implementation and configuration, has much in common with an MPLS Layer 2 VPN. In a VPLS topology, a packet originating within a customer's network is sent first to a customer edge (CE) device (for example, a router or Ethernet switch). It is then sent to a provider edge (PE) router within the service provider's network. The packet traverses the service provider's network over an MPLS label-switched path (LSP). It arrives at the egress PE router, which then forwards the traffic to the CE device at the destination customer site.

The difference is that for VPLS, packets can traverse the service provider's network in point-to-multipoint fashion, meaning that a packet originating from a CE device can be broadcast to all the PE routers participating in a VPLS routing instance. In contrast, a Layer 2 VPN forwards packets in point-to-point fashion only. The paths carrying VPLS traffic between each PE router participating in a routing instance are signaled using BGP.


NOTE: The RSVP automatic mesh feature with multiple RSVP neighbors on a single LAN is not supported on SRX Series devices because RSVP runs on WAN links in a service provider network. Most of these WAN interfaces are point-to-point and are rarely seen in LAN networks.

This topic contains the following sections:

Sample VPLS Topology

[Figure 119 on page 1547](#) shows a basic VPLS topology.

Figure 119: Basic VPLS Topology



In this sample, the PE routers use the same autonomous system (AS). Within the AS, routing information is communicated through an interior gateway protocol (IGP). Outside the AS, routing information is shared with other ASs through BGP. The PE routers must use the same signaling protocols to communicate.

VPLS on PE Routers

Within a VPLS configuration, a device running Junos OS can act as a PE router. Junos OS passes the VPLS traffic through the following ports and PIMs on the Juniper Networks device to CE routers in the VPLS network:

- Built-in Ethernet ports on front panel
- Gigabit Ethernet uPIMs
- Gigabit Ethernet ePIMs

- Fast Ethernet PIMs
- Fast Ethernet ePIMs

NOTE: Ports on uPIMs and ePIMs must be in routing mode before you can configure the corresponding interfaces for VPLS.


Because a VPLS carries Ethernet traffic across a service provider network, it must mimic an Ethernet network in some ways. When a PE router configured with a VPLS routing instance receives a packet from a CE device, it first determines whether it has the destination of the VPLS packet in the appropriate routing table. If it does, it forwards the packet to the appropriate PE router or CE device. If it does not, it broadcasts the packet to all other PE routers and CE devices that are members of that VPLS routing instance. In both cases, the CE device receiving the packet must be different from the one sending the packet.

When a PE router receives a packet from another PE router, it first determines whether it has the destination of the VPLS packet in the appropriate routing table. If it does, the PE router either forwards the packet or drops it depending on whether the destination is a local or remote CE device:

- If the destination is a local CE device, the PE router forwards the packet to it.
- If the destination is a remote CE device (connected to another PE router), the PE router discards the packet.

If the PE router cannot determine the destination of the VPLS packet, it floods the packet to all attached CE devices. [Figure 120 on page 1548](#) illustrates this process.

Figure 120: Flooding a Packet with an Unknown Destination



A VPLS interface can be directly connected to an Ethernet switch. Layer 2 information gathered by an Ethernet switch, for example, MAC addresses and interface ports, is included in the VPLS routing instance table.

An MPLS label-switched interface (LSI) label is used as the inner label for VPLS. This label maps to a VPLS routing instance on the ingress PE router. On the egress PE router, the LSI label is stripped and then mapped to a logical LSI interface. The Layer 2 Ethernet frame is then forwarded using the LSI interface to the correct VPLS routing instance.

One restriction on flooding behavior in VPLS is that traffic received from remote PE routers is never forwarded to other PE routers. This restriction helps prevent loops in the core network. However, if a CE Ethernet switch has two or more connections to the same PE router, you must enable the Spanning Tree Protocol (STP) on the CE switch to prevent loops.

NOTE: Under certain circumstances, VPLS PE routers might duplicate an Internet Control Message Protocol (ICMP) reply from a CE device when a PE router has to flood an ICMP request because the destination MAC address has not yet been learned. The duplicate ICMP reply can be triggered when a CE device with promiscuous mode enabled is connected to a PE router. The PE router automatically floods the promiscuous mode enabled CE device, which then returns the ICMP request to the VPLS PE routers. The VPLS PE routers consider the ICMP request to be new and flood the request again, creating a duplicate ping reply.

Using an Ethernet Switch as the VPLS CE Device

For VPLS configurations, the CE device does not necessarily need to be a router. You can link the PE routers directly to Ethernet switches. However, be aware of the following configuration issues:

- When you configure VPLS routing instances and establish two or more connections between a CE Ethernet switch and a PE router, you must enable the Spanning Tree Protocol (STP) on the switch to prevent loops.
- Junos OS allows standard bridge protocol data unit (BPDU) frames to pass through emulated Layer 2 connections, such as those configured with Layer 2 VPNs, Layer 2 circuits, and VPLS instances. However, CE Ethernet switches that generate proprietary BPDU frames might not be able to run STP across Juniper Networks routing platforms configured for these emulated Layer 2 connections.

VPLS Exceptions on SRX Series Devices

The VPLS implementation on SRX Series device is similar to VPLS implementations on M Series, T Series, and MX Series routers, with the following exceptions:

- SRX Series devices do not support aggregated Ethernet interfaces. Therefore, aggregated Ethernet interfaces between CE devices and PE routers, and aggregated Ethernet interfaces between PE devices and PE routers are not supported for VPLS routing instances on SRX Series devices.
- VPLS multihoming, which allows connecting a CE device to multiple PE routers to provide redundant connectivity, is not supported on SRX Series devices.
- SRX Series devices do not support BGP mesh groups.
- SRX Series devices support only the following encapsulation types on VPLS interfaces that face CE devices: extended VLAN VPLS, Ethernet VPLS, and VLAN VPLS. Ethernet VPLS over ATM LLC encapsulation is not supported.
- Virtual ports are generated dynamically on a Tunnel Services PIC on some Juniper Networks routing platforms. SRX Series devices do not support Tunnel Services modules or virtual ports.
- The VPLS implementation on SRX Series devices does not support dual-tagged frames. Therefore, VLAN rewrite operations are not supported on dual-tagged frames. VLAN rewrite operations such as pop-pop, pop-swap, push-push, swap-push, and swap-swap, which are supported on M Series and T Series routing platforms, are not supported on SRX Series devices.

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[Understanding VPLS Interfaces | 1561](#)

[Understanding VPLS Routing Instances | 1573](#)

[Understanding VPLS VLAN Encapsulation | 1631](#)

[Understanding VPLS VLAN Encapsulation on a Logical Interface | 1632](#)

[VPLS Configuration Overview | 1550](#)

VPLS Configuration Overview

To configure VPLS functionality, you must enable VPLS support on the provider edge (PE) routers. You must also configure PE routers to distribute routing information to the other PE routers in the VPLS and configure the circuits between the PE routers and the customer edge (CE) devices, as explained in the steps that follow.

NOTE: Many configuration procedures for VPLS are identical to the procedures for Layer 2 and Layer 3 VPNs.

To configure VPLS:

1. Determine which uPIM and ePIM ports correspond to the interfaces that will carry the VPLS traffic and enable routing mode on those ports.
2. Configure the interfaces that will carry the VPLS traffic between the PE router and CE devices. On the PE router interfaces that are facing the CE devices, specify a VPLS encapsulation type. The type of encapsulation depends on the interface type. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
3. Create a VPLS routing instance on each PE router that is participating in the VPLS. For each VPLS routing instance, specify which interfaces will carry the VPLS traffic between the PE and CE devices. On the CE device interface that faces the PE router, you must specify inet (for IPv4), and include the IP address. Additionally, each routing instance must have a unique route distinguisher associated with it. (VPN routing instances need a route distinguisher to help BGP identify overlapping network layer reachability information (NLRI) messages from different VPNs.) See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
4. Configure routing options on the PE router. See [“Example: Configuring Routing Options on the VPLS PE Router” on page 1630](#).
5. Configure MPLS LSPs between the PE routers. See [“Example: Configuring MPLS on the VPLS PE Router” on page 1585](#).
6. Configure RSVP on the PE routers. Enable RSVP for all connections that participate in the MPLS LSP. See [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#).
7. Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See [“Example: Configuring BGP on the VPLS PE Router” on page 1628](#).
8. Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
9. Configure VLAN encapsulation. See [“Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces” on page 1637](#), [“Example: Configuring VPLS VLAN Encapsulation” on page 1633](#), and [“Example: Configuring Extended VLAN VPLS Encapsulation” on page 1639](#).

RELATED DOCUMENTATION

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[MPLS Layer 2 VPN Configuration Overview | 1429](#)

[Example: Configuring MPLS on the VPLS PE Router | 1585](#)

[Example: Configuring RSVP on the VPLS PE Router | 1583](#)

[Example: Configuring BGP on the VPLS PE Router | 1628](#)

[Example: Configuring OSPF on the VPLS PE Router | 1581](#)

Migrating from FEC128 LDP-VPLS to EVPN Overview

IN THIS SECTION

- [Technology Overview and Benefits | 1552](#)
- [FEC128 LDP-VPLS to EVPN Migration | 1553](#)
- [Sample Configuration for LDP-VPLS to EVPN Migration | 1555](#)
- [Reverting to VPLS | 1559](#)
- [LDP-VPLS to EVPN Migration and Other Features | 1560](#)

For service providers with virtual private LAN service (VPLS) networks and Ethernet VPN (EVPN) networks, there is a need to interconnect these networks. Prior to Junos OS Release 17.3, a logical tunnel interface on the interconnection point of the VPLS and EVPN routing instances was used for this purpose. In this case, the provider edge (PE) devices in each network were unaware of the PE devices in the other technology network. Starting in Junos OS Release 17.3, a solution is introduced for enabling staged migration from FEC128 LDP-VPLS toward EVPN on a site-by-site basis for every VPN routing instance. In this solution, the PE devices running EVPN and VPLS for the same VPN routing instance and single-homed segments can coexist. During migration, there is minimal impact to the customer edge (CE) device-to-CE device traffic forwarding for affected customers.

The following sections describe the migration from LDP-VPLS to EVPN:

Technology Overview and Benefits

Virtual private LAN service (VPLS) is an Ethernet-based point-to-multipoint Layer 2 VPN. This technology allows you to connect geographically dispersed data center LANs to each other across an MPLS backbone while maintaining Layer 2 connectivity. The high availability features defined in VPLS standards (such as LER dual homing) and topology autodiscovery features using BGP signaling make VPLS scalable and easy

to deploy. Because VPLS uses MPLS as its core, it provides low latency variation and statistically bound low convergence times within the MPLS network.

Ethernet VPN (EVPN), on the other hand, is a combined Layer 2 and Layer 3 VPN solution that is more scalable, resilient, and efficient than current technologies. It provides several benefits including greater network efficiency, reliability, scalability, virtual machine (VM) mobility, and policy control for service providers and enterprises.

Although VPLS is a widely deployed Layer 2 VPN technology, service provider networks migrate to EVPN because of the scaling benefits and ease of deployment. Some of the benefits of EVPN include:

- Control plane traffic is distributed with BGP and the broadcast and multicast traffic is sent using a shared multicast tree or with ingress replication.
- Control plane learning is used for MAC and IP addresses instead of data plane learning. MAC address learning requires the flooding of unknown unicast and ARP frames, whereas IP address learning does not require any flooding.
- Route reflector is used to reduce a full mesh of BGP sessions among PE devices to a single BGP session between a PE device and the route reflector.
- Autodiscovery with BGP is used to discover PE devices participating in a given VPN, PE devices participating in a given redundancy group, tunnel encapsulation types, multicast tunnel type, and multicast members.
- All-active multihoming is used. This allows a given CE device to have multiple links to multiple PE devices, and traffic traversing to-and-from that CE device fully utilizes all of these links (Ethernet segment).
- When a link between a CE device and a PE device fails, the PE devices for that EVPN instance (EVI) are notified of the failure with the withdrawal of a single EVPN route. This allows those PE devices to remove the withdrawing PE device as a next hop for every MAC address associated with the failed link (mass withdrawal).

FEC128 LDP-VPLS to EVPN Migration

Some service providers want to preserve their investments in VPLS. This leads to the need to connect the old VPLS networks to new networks that run EVPN. For this purpose, logical tunnel interfaces on the interconnection point of the VPLS and EVPN routing instances were used. However, all the other PE devices belonged either to the VPLS network or to the EVPN network and were unaware of the other technology.

Starting in Junos OS Release 17.3, EVPN can be introduced into an existing VPLS network in a staged manner, with minimal impact to VPLS services. On a VPLS PE device, some customers can be moved to EVPN, while other customers continue to use VPLS pseudowires. Other PE devices can be entirely VPLS and switching customers on other PE devices to EVPN. This solution provides support for the seamless migration Internet draft (expires January ,2018), *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*.

The seamless migration from FEC128 LDP-VPLS to EVPN solution supports the following functionality:

- Allow for staged migration toward EVPN on a site-by-site basis per VPN instance. For instance, new EVPN sites to be provisioned on EVPN PE devices.
- Allow for the coexistence of PE devices running both EVPN and VPLS for the same VPN instance and single-homed segments.

In the LDP-VPLS to EVPN migration, the PE device where some customers have been migrated to EVPN while other customers are being served using VPLS is called a super PE device. As super PE devices discover other super PE devices within a routing instance, they use EVPN forwarding to communicate with other super PE devices and VPLS pseudowires to PE devices running VPLS. The PE device with no EVPN awareness, and running only VPLS for all the customers, is called a VPLS PE device.

The CE device connected to a super PE can reach CE devices connected to EVPN-only PE devices or VPLS-only PE devices, but CE devices connected to EVPN-only PE devices cannot reach CE devices connected to VPLS-only PE devices.

Because the migration from LDP-VPLS to EVPN is supported on a per-routing instance basis, and if the routing instance is serving multiple customers on a PE device, all are migrated together. EVPN is responsible for setting up data forwarding between the PE devices upgraded to EVPN, while VPLS continues to set up data forwarding to PE devices that run VPLS. There should be zero impact for customers that still use VPLS pseudowire on all the PE devices.

NOTE:

The following features are not supported with the LDP-VPLS to EVPN migration:

- Migration from FEC129 VPLS to EVPN.
- Migration from BGP-VPLS to EVPN.
- Migration of VPLS virtual switch to EVPN virtual switch.
- Migration of VPLS routing instance to EVPN virtual switch.
- Migration of VPLS routing instance or PBB-VPLS to PBB-EVPN.
- Seamless migration from EVPN back to VPLS.
- Enhancing EVPN to support the set of tools or statements and commands that VPLS supports.
- Active-active and active-standby multihoming. The migration to EVPN is supported only on single-homed deployments.
- Spanning all-active across EVPN and VPLS PE devices does not work, because the all-active multihoming feature is not supported on VPLS.
- Connecting EVPN-only PE devices with VPLS-only PE devices through super PE devices.
- IPv6, logical systems, multichassis support, and SNMP, because they are currently not supported on EVPN.

Sample Configuration for LDP-VPLS to EVPN Migration

IN THIS SECTION

- [LDP-VPLS Configuration | 1555](#)
- [EVPN Migration Configuration | 1557](#)

The following sections provide the sample configuration required for performing the LDP-VPLS to EVPN migration.

LDP-VPLS Configuration

A typical static LDP-VPLS routing instance configuration is as follows:

```
user@host# show routing-instance foo
```



```

instance-type vpls;
vlan-id 100; (not needed for VLAN bundle service)
interface ge-2/0/0.590;
interface ae500.590;
routing-interface irb.0;
forwarding-options {
    family vpls {
        filter {
            input UNKNOWN-UNICAST;
        }
    }
}
protocols {
    vpls {
        control-word;
        encapsulation-type ethernet-vlan;
        enable-mac-move-action;
        mac-table-size {
            100000;
            packet-action drop;
        }
        mac-table-aging-time ;
        interface-mac-limit {
            100000;
            packet-action drop;
        }
        no-tunnel-services; (use label-switched interfaces)
        vpls-id 245015;
        mtu 1552;
        ignore-mtu-mismatch;
        mac-flush {
            any-spoke;
        }
        no-vlan-id-validate;
        neighbor 192.168.252.64 {
            psn-tunnel-endpoint 10.0.0.31;
            pseudowire-status-tlv;
            revert-time 60;
            backup-neighbor 192.168.252.65 {
                psn-tunnel-endpoint 10.0.0.32;
                hot-standby;
            }
        }
    }
    mesh-group Spoke { (access label-switched interface toward spoke)

```



```

local-switching;
neighbor 192.168.252.66 {
    psn-tunnel-endpoint 10.0.0.41;
    pseudowire-status-tlv;
}
neighbor 192.168.252.67 {
    psn-tunnel-endpoint 10.0.0.42;
    pseudowire-status-tlv;
}
}
connectivity-type permanent;
}

```

```

user@host# show interfaces ge-2/0/0.590
encapsulation vlan-vpls;
vlan-id 590;
output-vlan-map {
    swap;
    tag-protocol-id 0x8100;
    inner-vlan-id 590;
}
family vpls {
    filter {
        input-list [ listA ];
        output-list listB;
    }
}

```

EVPN Migration Configuration

To perform the FEC128 LDP-VPLS to EVPN migration, do the following:

1. On the backup Routing Engine, load Junos OS Release 17.3R1.
2. Perform in-service software upgrade (ISSU) to acquire mastership. Ensure that the VPLS unified ISSU does not have any impact on the VPLS forwarding.
3. Identify routing instances (customers) that need to be migrated to EVPN.
4. Enable EVPN in a single routing instance.
 - Change routing instance type to **evpn**, and include the **evpn** statement at the **[edit routing-instances routing-instance-name protocols]** hierarchy level, and also include the **vpls** statement at the same hierarchy to support VPLS commands.

For example:

```
[edit routing-instances routing-instance-name]
instance-type evpn;
interface ge-2/0/0.590;
interface ae500.590;
routing-interface irb.0;
route-distinguisher 1.1.1.1:50; (add for LDP-VPLS)
vrf-target target:100:100; (add for LDP-VPLS)
forwarding-options {
  family vpls {
    filter {
      input UNKNOWN-UNICAST;
    }
  }
}
protocols {
  vpls { (supports all existing VPLS commands)
}
```

5. Enable family EVPN signaling in BGP.

For example:

```
protocols {
  bgp {
    local-as 102;
    group 2mx {
      type internal;
      local-address 81.1.1.1;
      family evpn {
        signaling;
      }
    }
    neighbor 81.2.2.2;
    neighbor 81.9.9.9;
  }
}
```

After the configuration for the EVPN migration is committed, the routing protocol process and the Layer 2 address learning process start building the EVPN state to reflect interfaces, bridge domains, peers and routes. The locally learnt MAC addresses are synchronized by the Layer 2 address learning process in the instance.vpls.0 to the routing protocol process. When a local MAC ages out in the instance.vpls.0, the routing protocol process is informed by the Layer 2 address learning process.

When an EVPN peer is learnt, the routing protocol process sends a new message to the Layer 2 address learning process to remove the peer's label-switched interface or virtual tunnel logical interface from the VE mesh group and disables MAC-learning on it. The EVPN IM next-hop is then added to the VE mesh group. The EVPN behavior in the routing protocol process of learning MAC addresses over BGP and informing Layer 2 address learning process of the MPLS next hop is maintained.

The VPLS statements and commands continue to apply to the VPLS pseudowires between the PE devices and the MAC addresses learnt over them. The EVPN statements and commands apply to PE devices running EVPN.

Reverting to VPLS

If the EVPN migration runs into issues, you can revert back to VPLS until the issue is understood. The routing instance is reverted from a super PE to a VPLS PE in a non-catastrophic manner by enabling the following configuration:

```
[edit routing-instances routing-instance-name]
user@host# set instance-type vpls
user@host# delete protocols evpn
user@host# delete route-distinguisher (if running LDP-VPLS)
user@host# delete vrf-target (if running LDP-VPLS)
```

On reverting the EVPN migration to VPLS, the following happens:

1. The EVPN state information is deleted.
2. There is a trigger for withdrawal of EVPN control plane routes.
3. The routing protocol process sends a new message to the Layer 2 address learning process with the label-switched interface or the virtual tunnel logical interface for the routing instance and peer.
4. The label-switched or virtual tunnel interface adds the new message to the flood group and MAC learning is enabled.
5. The egress IM next hop is deleted by the routing protocols process, prompting the Layer 2 address learning process to remove it from the flood group.
6. Remote MAC addresses are learned again over the label-switched interface or virtual tunnel logical interface.

LDP-VPLS to EVPN Migration and Other Features

Table 46 on page 1560 describes the functionality of some of the related features, such as multihoming and integrated routing and bridging (IRB) with the LDP-VPLS to EVPN migration.

Table 46: EVPN Migration and Other Features Support

Feature	Supported Functionality in EVPN Migration
MAC move	<p>MAC moves are supported between VPLS-only PE device and super PE devices.</p> <p>When a MAC address moves from a VPLS-only PE device to a super PE device, it is learned over BGP, and the routing protocol process informs the Layer 2 address learning process of the EVPN next hop to be updated in the foo.vpls.0 routing table.</p> <p>When a MAC address moves from a super PE device to a VPLS-only PE device, it is learned in the Packet Forwarding Engine on the label-switched interface or virtual tunnel interface. The Layer 3 address learning process updates it to VPLS or the label-switched interface next hop.</p> <p>When the type 2 route is withdrawn by EVPN BGP, the MAC address is not deleted from the forwarding table, so there is no loss of data.</p> <p>The forwarding MAC table is shared by VPLS and EVPN. Some attributes, such as mac-table-size and mac-table-aging-time could be configured under both EVPN and VPLS. When there is a conflict, the values under EVPN take precedence.</p>
IRB	<p>No changes needed in IRB.</p> <p>On a super PE device, EVPN populates the /32 host routes learned over MAC+IP type 2 routes from EVPN peers in a Layer 3 virtual routing and forwarding, while VPLS IRB forwarding using subnet routes works on sites still running VPLS.</p>

Table 46: EVPN Migration and Other Features Support *(continued)*

Feature	Supported Functionality in EVPN Migration
Hierarchical VPLS	<p>In an H-VPLS network with hub-and-spoke PE devices, when the hub PE device is migrated to EVPN, local MAC addresses learned over the access label-switched or virtual tunnel interface need to be advertised to BGP, so that the other EVPN-only PE devices or super PE devices can reach them.</p> <p>Take the following into consideration when migrating an H-VPLS network to EVPN:</p> <ul style="list-style-type: none">• Hubs typically have local switching enabled as interspoke traffic is forwarded through the hub. If spokes alone are migrated to EVPN and spokes have Layer 3 or MPLS reachability to each other, the label-switched or virtual tunnel interface to the hub and EVPN next hop (remote spoke) is present in the VPLS edge (VE) floodgroup. This results in two copies of broadcast, unknown unicast, and multicast (BUM) traffic received by the remote spoke. One option to avoid this behavior is to migrate the hubs to EVPN too.• EVPN is not aware of hierarchy. All peers are considered core-facing. Once hubs and spokes are migrated to EVPN, split horizon prevents the BUM traffic from being forwarded to other core-facing PE devices.
ESI configuration	Ethernet segment identifier (ESI) is configured at the physical interface or port level.

RELATED DOCUMENTATION

| [EVPN Overview](#)

Understanding VPLS Interfaces

IN THIS SECTION

- [Interface Name | 1562](#)
- [Encapsulation Type | 1562](#)
- [Flexible VLAN Tagging | 1562](#)
- [VLAN Rewrite | 1563](#)

For each VPLS routing instance on a PE router, you specify which interfaces are to be used to carry VPLS traffic between the PE and CE devices.

This topic contains the following sections:

Interface Name

Specify both the physical and logical portions of the interface name, in the following format:

physical.logical

For example, in ge-1/2/1.2, ge-1/0/1 is the physical portion of the interface name and 2 is the logical portion. If you do not specify the logical portion of the interface name, 0 is set by default. A logical interface can be associated with only one routing instance.

Encapsulation Type

The physical link-layer encapsulation type for a VPLS interface can be one of the following:

- **ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol Identifier (TPID) values.
- **extended-vlan-vpls**—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. All VLAN IDs from 1 through 1023 are valid for VPLS VLANs on Fast Ethernet interfaces, and all VLAN IDs from 1 through 4094 are valid for VPLS VLANs on Gigabit Ethernet interfaces.
- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. You must configure this encapsulation type on both the physical interface and the logical interface. VLAN IDs 1 through 511 are reserved for normal Ethernet VLANs, IDs 512 through 1023 are reserved for VPLS VLANs on Fast Ethernet interfaces, and IDs 512 through 4094 are reserved for VPLS VLANs on Gigabit Ethernet interfaces.
- **flexible-ethernet-services**—Use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. This encapsulation type allows you to configure any combination of route, TCC, CCC, and VPLS encapsulations on a single physical port. Aggregated Ethernet bundles cannot use this encapsulation type.

For flexible Ethernet services encapsulation, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

Flexible VLAN Tagging

For untagged packets to be accepted on an 802.1Q VLAN-tagged port, specify the native VLAN ID with the flexible VLAN tagging option. (No other flexible VLAN tagging features are supported.)

VLAN Rewrite

You can rewrite VLAN tags on VPLS interfaces. Rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between CE devices that share a VLAN ID.

You can configure rewrite operations to stack (push), remove (pop), or rewrite (swap) tags on single-tagged frames. If a port is not configured for VLAN tagging, rewrite operations are not supported on any logical interface on that port.

You can configure the following VLAN rewrite operations:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- **push**—Add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.
- **swap**—Replace the VLAN tag at the top of the VLAN tag stack with a user-specified VLAN tag value.

You perform VLAN rewrite operations by applying input and output VLAN maps at the ingress and egress, respectively, of the interface. For incoming frames, use the `input-vlan-map`; for outgoing frames, use the `output-vlan-map`.

The VPLS implementation on SRX Series devices does not support dual-tagged frames. Therefore, VLAN rewrite operations are not supported on dual-tagged frames. VLAN rewrite operations such as pop-pop, pop-swap, push-push, swap-push, and swap-swap, which are supported on M Series and T Series routing platforms, are not supported on SRX Series devices.

RELATED DOCUMENTATION

[Example: Configuring Routing Interfaces on the VPLS PE Router | 1564](#)

[Example: Configuring the Interface to the VPLS CE Device | 1565](#)

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

[Understanding VPLS VLAN Encapsulation | 1631](#)

Example: Configuring Routing Interfaces on the VPLS PE Router

IN THIS SECTION

- [Requirements | 1564](#)
- [Overview | 1564](#)
- [Configuration | 1564](#)
- [Verification | 1565](#)

This example shows how to configure routing interfaces on the VPLS PE router.

Requirements

Before you begin, see *Understanding Selective Stateless Packet-Based Services* .

Overview

In this example, you configure the PE1 router loopback interface and the interface to the PE2 router ge-2/0/1.

Configuration

Step-by-Step Procedure

To configure the routing interface on the VPLS PE router:

1. Configure the loopback interface.

```
[edit]
user@host# set interfaces lo0 unit 0 family inet address 10.255.7.168/32 primary
```

2. Configure the IP address on the MPLS core interface.

```
[edit]
user@host# set interfaces ge-3/0/2 unit 0 family inet address 100.1.1.1/30
```

3. Configure the MPLS family.


```
[edit]
user@host# set interfaces ge-3/0/2 unit 0 family mpls
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

RELATED DOCUMENTATION

Interfaces User Guide for Security Devices

[VPLS Configuration Overview | 1550](#)

[Understanding VPLS Interfaces | 1561](#)

Example: Configuring the Interface to the VPLS CE Device

IN THIS SECTION

- [Requirements | 1565](#)
- [Overview | 1566](#)
- [Configuration | 1566](#)
- [Verification | 1566](#)

This example shows how to configure the router interface that is connected to the CE device to include VPLS encapsulation.

Requirements

Before you begin, see *Understanding Selective Stateless Packet-Based Services* .

Overview

In this example, you configure the router interface ge-1/2/1 that is connected to the CE device to include VPLS encapsulation.

Configuration

Step-by-Step Procedure

To configure the interface to the VPLS CE device:

1. Configure VPLS encapsulation for the interface facing the CE router.

```
[edit]
user@host# set interfaces ge-1/2/1 encapsulation ethernet-vpls
```

2. Configure the interface for the VPLS family group.

```
[edit]
user@host# set interfaces ge-1/2/1 unit 0 family vpls
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces ge-1/2/1** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[Understanding VPLS Interfaces | 1561](#)

VPLS Filters and Policers Overview

This feature permits users to configure both firewall filters and policers for virtual private LAN service (VPLS). Firewall filters enable you to filter packets based on their components and perform an action on packets that match the filter. Policers enable you to limit the amount of traffic that passes into or out of an interface.

This feature can be enabled by configuring VPLS filters, policers, and accounting through various CLI commands. VPLS filters and policers act on a Layer 2 frame that includes the media access control (MAC) header (after any VLAN rewrite or other rules are applied), but that does not include the cyclical redundancy check (CRC) field.

NOTE: You can apply VPLS filters and policers on the PE routers only to customer-facing (PE-CE) interfaces.

RELATED DOCUMENTATION

[Example: Configuring VPLS Policers | 1571](#)

[Example: Configuring VPLS Filters | 1567](#)

Example: Configuring VPLS Filters

IN THIS SECTION

- [Requirements | 1568](#)
- [Overview | 1568](#)
- [Configuration | 1568](#)
- [Verification | 1570](#)

This example shows how to configure VPLS filters.

Requirements

Before you begin:

- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE on the PE routers. See [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#).

Overview

This example describes how to configure filtering and accounting for VPLS.



CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your device to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the device performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device.

Configuration

CLI Quick Configuration

To quickly configure VPLS filters, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set firewall family vpls filter blue term term1 from interface ge-3/0/0.512
set firewall family vpls filter blue term term1 from interface fe-5/0/0.512
set firewall family vpls filter blue term term1 then count count1
set firewall family vpls filter blue accounting-profile fw_profile
set accounting-options file fw_acc size 500k
set accounting-options file fw_acc transfer-interval 5
```



```

set accounting-options filter-profile fw_profile file fw_acc
set accounting-options filter-profile fw_profile interval 1
set accounting-options filter-profile fw_profile counters count1
set interfaces ge-0/0/1 unit 512 family vpls filter input blue

```

Step-by-Step Procedure

To configure filters for VPLS:

1. Configure a filter with a GE interface as the match condition and count as the action.

```

[edit ]
user@host# set firewall family vpls filter blue term term1 from interface ge-3/0/0.512

```

2. Configure a filter with an FE interface as the match condition and count as the action.

```

[edit ]
user@host# set firewall family vpls filter blue term term1 from interface fe-5/0/0.512

```

3. Configure the count.

```

[edit ]
user@host# set firewall family vpls filter blue term term1 then count count1

```

4. Configure the accounting profile to refer it to the counter.

```

[edit ]
user@host# set firewall family vpls filter blue accounting-profile fw_profile

```

5. Configure the account file size.

```

[edit ]
user@host# set accounting-options file fw_acc size 500k

```

6. Configure the account transfer interval.

```

[edit ]
user@host# set accounting-options file fw_acc transfer-interval 5

```


7. Configure the filter for the accounting profile.

```
[edit ]
user@host# set accounting-options filter-profile fw_profile file fw_acc
```

8. Configure the filter for the interval.

```
[edit ]
user@host# set accounting-options filter-profile fw_profile interval 1
```

9. Configure the counter.

```
[edit ]
user@host# set accounting-options filter-profile fw_profile counters count1
```

10. Apply the filter to the interface.

```
[edit ]
user@host# set interfaces ge-0/0/1 unit 512 family vpls filter input blue
```

11. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show firewall** and **show accounting records** commands.

RELATED DOCUMENTATION

[VPLS Filters and Policers Overview | 1567](#)

[VPLS Configuration Overview | 1550](#)

[Example: Configuring VPLS Policers | 1571](#)

Example: Configuring VPLS Policers

IN THIS SECTION

- [Requirements | 1571](#)
- [Overview | 1571](#)
- [Configuration | 1572](#)
- [Verification | 1573](#)

This example shows how to configure VPLS policers.

Requirements

Before you begin:

- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE on the PE routers. See [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#).

Overview

This example describes how to configure policing and apply it on the interface for VPLS.



CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your device to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the device performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the device.

Configuration

CLI Quick Configuration

To quickly configure VPLS policers, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set firewall policer police2 if-exceeding bandwidth-percent 10
set firewall policer police2 if-exceeding burst-size-limit 1500
set firewall policer police2 then discard
set interfaces ge-0/0/1 unit 512 family vpls policer input police2
```

Step-by-Step Procedure

To configure filters for VPLS:

1. Configure bandwidth percentage.

```
[edit ]
user@host# set firewall policer police2 if-exceeding bandwidth-percent 10
```

2. Configure the burst size limit.

```
[edit ]
user@host# set firewall policer police2 if-exceeding burst-size-limit 1500
```

3. Configure the terminal action on the packet.

```
[edit ]
user@host# set firewall policer police2 then discard
```

4. Apply the policer to the interface.

```
[edit ]
user@host# set interfaces ge-0/0/1 unit 512 family vpls policer input police2
```

5. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```


Verification

To verify the configuration is working properly, enter the **show firewall** command.

RELATED DOCUMENTATION

[VPLS Filters and Policers Overview | 1567](#)

[VPLS Configuration Overview | 1550](#)

[Example: Configuring VPLS Filters | 1567](#)

Understanding VPLS Routing Instances

IN THIS SECTION

- [BGP Signaling | 1574](#)
- [VPLS Routing Table | 1575](#)
- [Trace Options | 1576](#)

To configure VPLS functionality, you must enable VPLS support on the PE router. You must also configure PE routers to distribute routing information to the other PE routers in the VPLS and configure the circuits between the PE routers and the CE devices.

You create a VPLS routing instance on each PE router that is participating in the VPLS. The routing instance has the same name on each PE router. To configure the VPLS routing instance, you specify the following:

- **Route distinguisher**—Helps BGP distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPLS instances. Each routing instance that you configure on a PE router must have a unique route distinguisher.
- **Route target**—Defines which route is part of a VPLS. A unique route target helps distinguish between different VPLS services on the same router.
- **Site name**—Provides unique name for the VPLS site.
- **Site identifier**—Provides unique numerical identifier for the VPLS site.

- Site range—Specifies total number of sites in the VPLS. The site range must be greater than the site identifier.
- Interface to the CE router—Specifies the physical interface to the CE router that carries VPLS traffic. The interface must be configured for a VPLS encapsulation type.

NOTE: In addition to the VPLS routing instance, you must configure MPLS label-switched paths (LSPs) between the PE routers, internal BGP (IBGP) sessions between the PE routers, and an interior gateway protocol (IGP) on the PE routers.



CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your router to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the router performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the router.

This topic contains the following sections:

BGP Signaling

BGP is used to signal the paths between each of the PE routers participating in the VPLS routing instance. These paths carry VPLS traffic across the service provider's network between the VPLS sites.

NOTE: LDP signaling is not supported for the VPLS routing instance.

To configure BGP signaling, you specify the following:

- VPLS site name and site identifier—When you configure BGP signaling for the VPLS routing instance, you must specify each VPLS site that has a connection to the router. For each VPLS site, you must configure a site name and site identifier (a numerical identifier between 1 to 65,534 that uniquely identifies the VPLS site).
- Site range—When you enable BGP signaling for the VPLS routing instance, you need to configure a site range. The site range specifies the total number of sites in the VPLS.

NOTE: The site range value must be greater than the largest site identifier.

- **Site preference**—You can specify the preference value advertised for a particular VPLS site. The site preference value is encoded in the BGP local preference attribute. When a PE router receives multiple advertisements with the same VPLS edge (VE) device identifier, the advertisement with the highest local preference value is preferred.

VPLS Routing Table

The VPLS routing table contains MAC addresses and interface information for both physical and virtual ports. You can configure the following characteristics for the table:

- **Table size**—You can modify the size of the VPLS MAC address table. The default table size is 512 MAC addresses; the minimum is 16 addresses, and the maximum is 65,536 addresses.

If the MAC table limit is reached, new MAC addresses can no longer be added to the table. Eventually the oldest MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added. However, as long as the table is full, new MAC addresses are dropped.

The interfaces affected include all of the interfaces within the VPLS routing instance, including the local interfaces and the LSI interfaces.

- **Timeout interval**—You can modify the timeout interval for the VPLS table. The default timeout interval is 300 seconds; the minimum is 10 seconds, and the maximum is 1,000,000 seconds. We recommend you configure longer values for small, stable VPLS networks and shorter values for large, dynamic VPLS networks. If the VPLS table does not receive any updates during the timeout interval, the router waits one additional interval before automatically clearing the MAC address entries from the VPLS table.
- **Number of addresses learned from an interface**—You can configure a limit on the number of MAC addresses learned by a VPLS routing instance by setting the MAC table size. The default is 512 addresses; the minimum is 16, and the maximum is 65,536 addresses. If the MAC table limit is reached, new MAC addresses can no longer be added to the table. Eventually the oldest MAC addresses are removed from the MAC address table automatically. This frees space in the table, allowing new entries to be added. However, as long as the table is full, new MAC addresses are dropped.

Because this limit applies to each VPLS routing instance, the MAC addresses of a single interface can consume all the available space in the table, preventing the routing instance from acquiring addresses from other interfaces. You can limit the number of MAC addresses learned from all interfaces configured for a VPLS routing instance, as well as limit the number of MAC addresses learned from a specific interface.

The MAC limit configured for an individual interface overrides the limit configured for all interfaces for the VPLS routing instance. Also, the table limit can override the limits configured for the interfaces.

The MAC address limit applies only to interfaces to CE devices.

Trace Options

The following trace flags display operations associated with VPLS:

- **all**—All VPLS tracing options
- **connections**—VPLS connections (events and state changes)
- **error**—Error conditions
- **nlri**—VPLS advertisements received or sent using BGP
- **route**—Trace-routing information
- **topology**—VPLS topology changes caused by reconsideration or advertisements received from other PE routers using BGP

RELATED DOCUMENTATION

[Example: Configuring the VPLS Routing Instance | 1576](#)

[Example: Configuring Routing Options on the VPLS PE Router | 1630](#)

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

[Understanding VPLS VLAN Encapsulation | 1631](#)

Example: Configuring the VPLS Routing Instance

IN THIS SECTION

- [Requirements | 1577](#)
- [Overview | 1577](#)
- [Configuration | 1577](#)
- [Verification | 1579](#)

This example shows how to create a VPLS routing instance on each PE router that is participating in the VPLS.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).

Overview

This example describes how to create a VPLS routing instance; configure VPLS site identifier, site range, no tunnel services option, route distinguisher, and route target for the VPLS routing instance; and specify the VPLS interface to the CE router.

NOTE: You must specify no tunnel services in the VPLS routing instance configuration, because SRX Series devices do not support tunnel serial PICs.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set routing-instances green instance-type vpls
set routing-instances green protocols vpls site-range 10 site R3 site-identifier 2
set routing-instances green protocols vpls no-tunnel-services
set routing-instances green route-distinguisher 10.255.7.1:1
set routing-instances green vrf-target target:11111:1
set routing-instances green instance-type vpls interface ge-1/2/1.0
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a VPLS routing instance:

1. Configure the routing instance of type VPLS.


```
[edit]
user@host# edit routing-instances green
```

2. Enable the VPLS instance type.

```
[edit routing-instances green]
user@host# set instance-type vpls
```

3. Configure the VPLS site identifier and range for the VPLS routing instance.

```
[edit routing-instances green protocols vpls]
user@host# set site-range 10 site R3 site-identifier 2
```

4. Configure the no-tunnel-services option for the VPLS routing instance.

```
[edit routing-instances green protocols vpls]
user@host# set no-tunnel-services
```

5. Configure the route distinguisher.

```
[edit routing-instances green]
user@host# set route-distinguisher 10.255.7.1:1
```

6. Configure the route target.

```
[edit routing-instances green]
user@host# set vrf-target target:11111:1
```

7. Specify the VPLS interface to the CE router.

```
[edit routing-instances green]
user@host# set instance-type vpls interface ge-1/2/1.0
```

Results

From configuration mode, confirm your configuration by entering the **show routing-instances green** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.


```
[edit]
user@host# show routing-instances green
instance-type vpls;
interface ge-1/2/1.0;
route-distinguisher 10.255.7.1:1;
vrf-target target:11111:1;
protocols {
  vpls {
    site-range 10;
    no-tunnel-services;
    site R3 {
      site-identifier 2;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying VPLS Routing Instance Is Configured | 1579](#)
- [Verifying VPLS Routing Attributes Are Configured | 1579](#)

Confirm that the configuration is working properly.

Verifying VPLS Routing Instance Is Configured

Purpose

Verify that the VPLS routing instance is configured.

Action

From operational mode, enter the **show routing-instances** command.

Verifying VPLS Routing Attributes Are Configured

Purpose

Verify that attributes such as VPLS site identifier, site range, no tunnel services option, route distinguisher, and route target for the VPLS routing instance are configured.

Action

From operational mode, enter the **show routing-instances green protocols vpls** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[Understanding VPLS Routing Instances | 1573](#)

Example: Configuring Automatic Site Identifiers for VPLS

This example shows how to configure automatic site identifiers for VPLS sites.

Requirements

Before you begin, see information on selective stateless packet-based services in *Interfaces User Guide for Security Devices*.

Overview

When you enable automatic site identifiers, the Junos OS automatically assigns site identifiers to VPLS sites. In this example, you configure a routing instance called `vpls` instance and enable automatic site identifiers for VPLS.

NOTE: Site identifiers for VPLS sites can be different for different routing instances.

Configuration

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure automatic site identifiers:

1. Configure the routing instance of type VPLS.

[edit]


```
user@host#set routing-instances vpls-instance
```

2. Enable automatic site identifiers.

```
[edit routing-instances vpls-instance]
user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id collision-detect-time 10
user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id new-site-wait-time 20
user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id reclaim-wait-time minimum
5 maximum 20
user@host#set protocols vpls no-tunnel-services site site10 automatic-site-id startup-wait-time 5
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show vpls connections** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring OSPF on the VPLS PE Router

IN THIS SECTION

- [Requirements | 1582](#)
- [Overview | 1582](#)
- [Configuration | 1582](#)
- [Verification | 1583](#)

This example shows how to configure OSPF on the VPLS PE router.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).

Overview

The PE routers exchange routing information using an IGP such as OSPF. In this example, you configure OSPF area 0.0.0.0 on the VPLS PE router and traffic engineering for OSPF.

Configuration

Step-by-Step Procedure

To configure OSPF on the VPLS PE router:

1. Configure the OSPF area on the VPLS PE router.

```
[edit]
user@host# set protocols ospf area 0.0.0.0 interface t1-1/0/1.0
user@host# set protocols ospf area 0.0.0.0 interface lo0.0
```

2. Configure traffic engineering for OSPF.

```
[edit]
user@host# set protocols ospf traffic-engineering
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```


Verification

To verify the configuration is working properly, enter the **show protocols** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring RSVP on the VPLS PE Router

IN THIS SECTION

- [Requirements | 1583](#)
- [Overview | 1584](#)
- [Configuration | 1584](#)
- [Verification | 1584](#)

This example shows how to configure RSVP on the VPLS PE router.

Requirements

Before you begin:

- Before you begin, see *Understanding Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).

Overview

This example describes how to enable RSVP for all connections that participate in the LSP on the PE1 router.

Configuration

Step-by-Step Procedure

To configure RSVP on the VPLS PE router:

1. Configure the interface to the PE2 router for RSVP.

```
[edit ]
user@host# set protocols rsvp interface t1-1/0/1.0
```

2. Configure the loopback interface for RSVP.

```
[edit ]
user@host# set protocols rsvp interface lo0.0
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show protocols** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring MPLS on the VPLS PE Router

IN THIS SECTION

- [Requirements | 1585](#)
- [Overview | 1585](#)
- [Configuration | 1586](#)
- [Verification | 1586](#)

This example shows how to configure MPLS on the VPLS PE router.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE on the PE routers. See [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#).

Overview

This example shows you how to configure MPLS on the PE1 router to advertise the Layer 2 VPN interface that communicates with the PE2 router.



CAUTION: MPLS is disabled by default on SRX Series devices. You must explicitly configure your router to allow MPLS traffic. However, when MPLS is enabled, all flow-based security features are deactivated and the router performs packet-based processing. Flow-based services such as security policies, zones, NAT, ALGs, chassis clustering, screens, firewall authentication, and IPsec VPNs are unavailable on the router.

Configuration

Step-by-Step Procedure

To configure MPLS on the VPLS PE router:

1. Configure the interface to the PE2 router for MPLS.

```
[edit ]
user@host# set protocols mpls interface t1-1/0/1.0
```

2. Configure the loopback for MPLS.

```
[edit ]
user@host# set protocols mpls interface lo0.0
```

3. Configure the path to destination 10.255.7.164.

```
[edit ]
user@host# set protocols mpls label-switched-path chelsea-sagar to 10.255.7.164
```

4. If you are done configuring the device, commit the configuration.

```
[edit ]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show mpls** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring LDP on the VPLS PE Router

IN THIS SECTION

- [Requirements | 1587](#)
- [Overview | 1587](#)
- [Configuration | 1587](#)
- [Verification | 1588](#)

This example shows how to configure LDP on the VPLS PE router.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).

Overview

This example describes how to enable LDP for all connections that participate in the LSP on the PE1 router.

Configuration

Step-by-Step Procedure

To configure LDP on the VPLS PE router:

1. Configure the interface to the PE2 router for LDP.

```
[edit ]
user@host# set protocols ldp interface ge-3/0/2
```

2. Configure the loopback interface for LDP.

```
[edit ]
user@host# set protocols ldp interface lo0
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show protocols** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring VPLS over GRE with IPsec VPNs

IN THIS SECTION

- [Requirements | 1589](#)
- [Overview | 1589](#)
- [Configuration | 1594](#)
- [Verification | 1609](#)

This example demonstrates a network scenario consisting of a central office and one branch office that will use VPLS, MPLS, GRE, and IPsec to create secure Ethernet connectivity over a Layer 3 network. This configuration can be expanded to add many other branch sites.

Requirements

Before you begin:

- Ensure that a layer 3 network is in place for all branch offices and that there is an ingress (head-end) device at the central office configured to terminate the VPNs from each branch office.
- Obtain IDP licenses for each SRX Series device. IDP is used to reassemble GRE packets that might become fragmented.

Overview

Junos OS can selectively choose whether traffic is processed by the flow engine or packet engine using the selective stateless packet-based feature. This feature allows you to combine flow and packet-based services in a single device. In this example, we describe a deployment scenario that uses this feature to deploy large-scale VPLS over GRE. This enables SRX devices to securely transport Ethernet traffic over Layer 3 networks when used in conjunction with IPsec.

In this scenario you configure a central office ingress (head-end) using an SRX650 device and one branch office using an SRX240 device. This setup is accomplished by carrying MPLS pseudowires over GRE, which in turn, is encapsulated in IPsec in order to guarantee data integrity and confidentiality. By default, SRX Series devices use secure flow forwarding. Because VPLS services are provided in packet-mode only, the configuration requires the GRE tunnel to be terminated in a packet-mode routing instance (the default routing instance).

NOTE: You can also use an MX Series device as the ingress (head-end) device, which is mentioned later in this topic.

To better understand this configuration, we will discuss two scenarios. The first scenario uses pseudowires to allow the creation of point-to-point circuits between two endpoints carried over the MPLS network. If we leave the signaling protocols aside (that is, there are a few ways to provision the pseudowires), these connections are just point-to-point connections. Using this approach provides an end-to-end wire between sites. This is beneficial from a traffic processing point of view because the gateways do not need to do MAC address learning, they simply forward anything they receive to the pseudowire. Because of this, it may be difficult to deploy this setup when trying to provide connectivity to multiple branch offices.

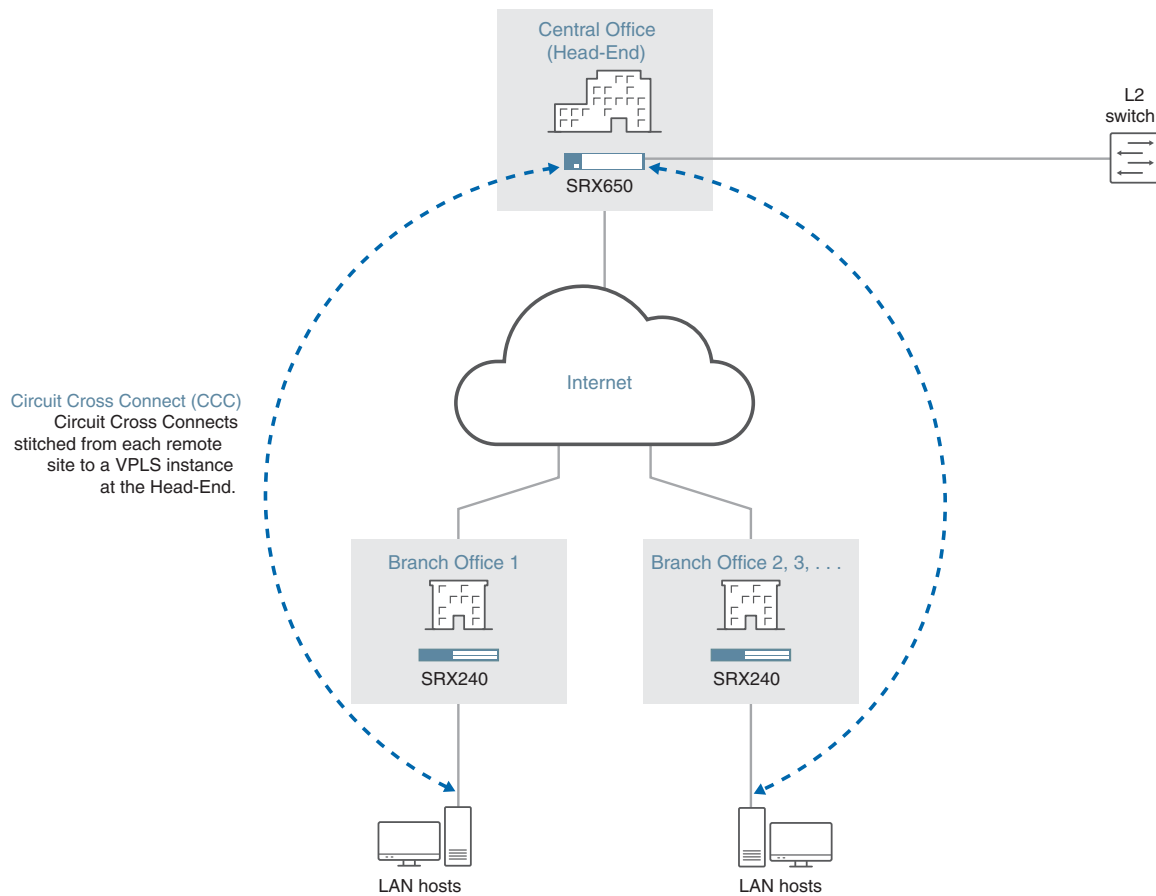
The second scenario could use VPLS to provide a Layer 2 network abstraction. With VPLS, endpoints are expected to negotiate LSPs and pseudowires with every other endpoint (that is, they are fully meshed).

When a node receives an Ethernet frame from one of its LAN interfaces the source MAC address is learned, if it's not already known, and flooded using every pseudowire connecting to all other branch nodes. However, if the destination has been previously learned, then the frame is sent to the appropriate destination. When an Ethernet frame is received through one of the pseudowires (that is, from the MPLS network), source MAC address learning is performed. The next time a frame is sent to that MAC it does not need to be flooded and the frame is flooded to every single LAN interface in the node, but not over the pseudowires. In other words, the network acts as a distributed Layer 2 switch providing any-to-any Ethernet connectivity between the devices connected to the different nodes in the network.

While the advantages of this second scenario is evident (any-to-any connectivity, automated provisioning, and simple abstraction), it comes at the cost of complexity. Every PE node has to perform Layer 2 learning and flooding of traffic, which can cause problems when either multiple broadcast/multicast or frames to unknown MAC addresses are used. As an example, if you had a topology with a thousand branch offices, each office that receives a broadcast packet must replicate it 999 times, encapsulate each copy in GRE and IPsec and forward the resulting traffic. Additionally, because each node performs Layer 2 learning, there are limitations in the maximum number of MAC addresses that each node can learn, limiting the total number of nodes in the domain.

In this example, we use a hybrid approach to these two scenarios. We use a circuit cross connect (CCC) at each branch office stitched to a VPLS instance at central office (ingress). This solution makes sense if most of the traffic flows from the branch offices to central office, and the branch-to-branch office traffic is always forwarded through the hub. The use of CCCs at branch offices combined with VPLS stitching at the central office provides a scalable way to deploy large hub-and-spoke topologies where Ethernet must be transported over an IP network (with or without encryption). At the expense of configuration complexity, it is possible to use SRX Series devices to terminate such connections, providing a scalable and cost-effective way to deploy small-to-large networks where Ethernet traffic is carried transparently using lower cost IP connections. [Figure 121 on page 1591](#) shows this topology.

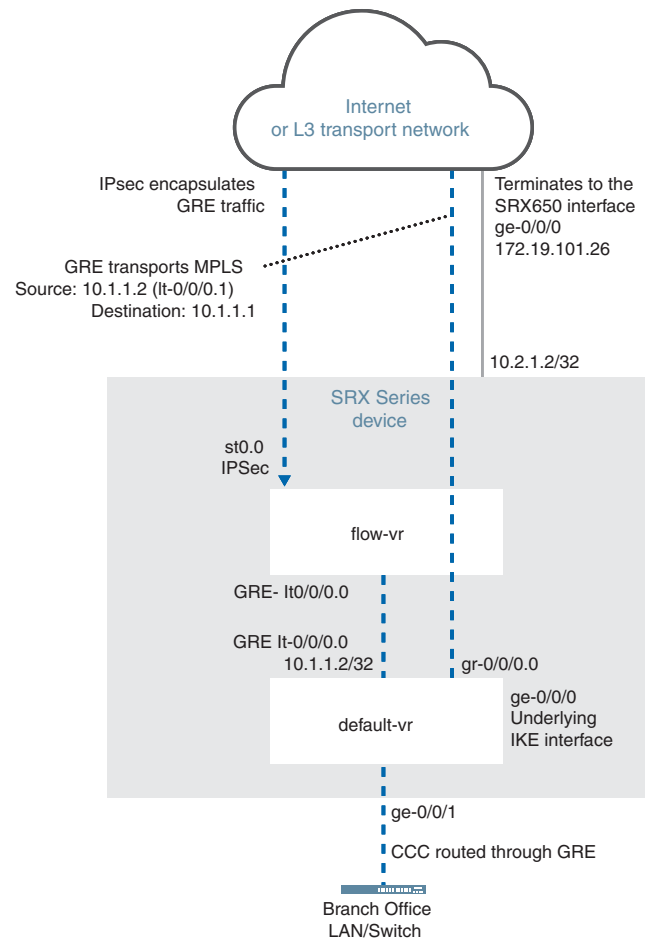
Figure 121: VPLS Deployment Scenario



9031178

In this deployment, VPLS services are provided only in packet mode and must be configured in the default routing instance. Unfortunately, IPsec is only provided in flow mode. Hence, a flow-mode routing-instance is used that provides both GRE reassembly and IPsec termination. While the GRE termination is done in the default routing instance, a flow-mode routing instance is connected between the default routing instance and the Internet (or whatever Layer 3 network is used as a transport), and it terminates the IPsec tunnel towards the ingress device. Because it is likely that a single public IP address is available, the Internet-facing Interface is connected to the default routing instance and is used to terminate IKE; however, the tunnel interface (st0) is bound to the flow-mode routing instance. See [Figure 122 on page 1592](#).

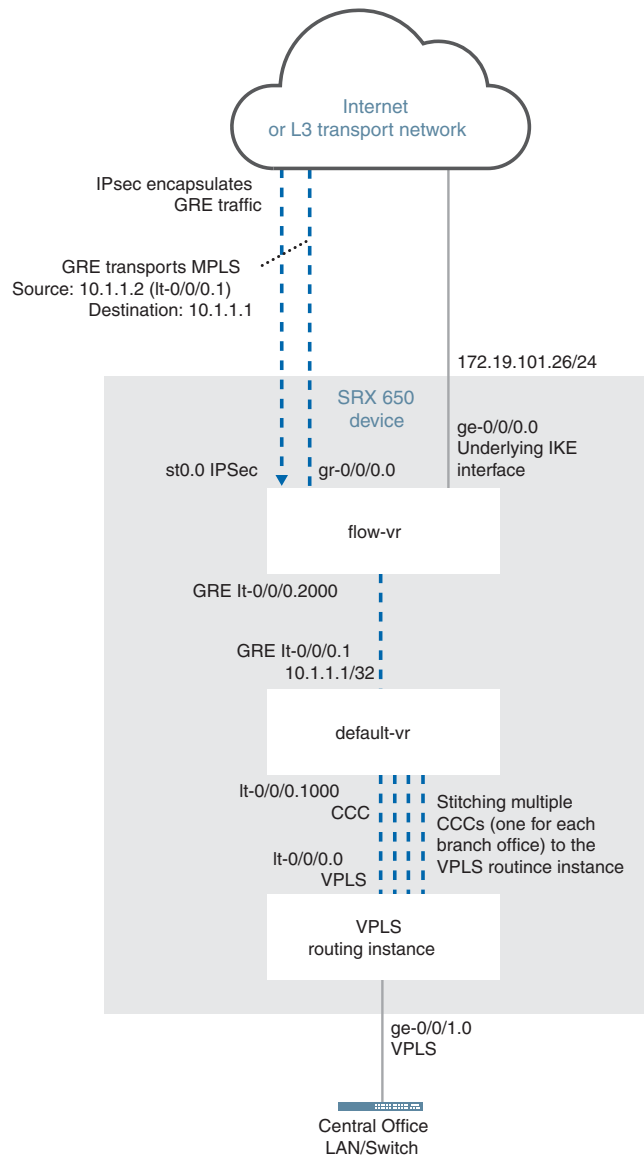
Figure 122: Branch Office Circuit Cross Connect Termination



When configuring the central office SRX650, the first thing you do is terminate the IPsec tunnels, GRE, and CCC connections. Because a SRX Series device is used as the ingress (head-end), the configuration to terminate the CCC circuits is identical to the one used at each branch office, with the exception that instead of one tunnel, multiple tunnels (and pseudowires) are terminated.

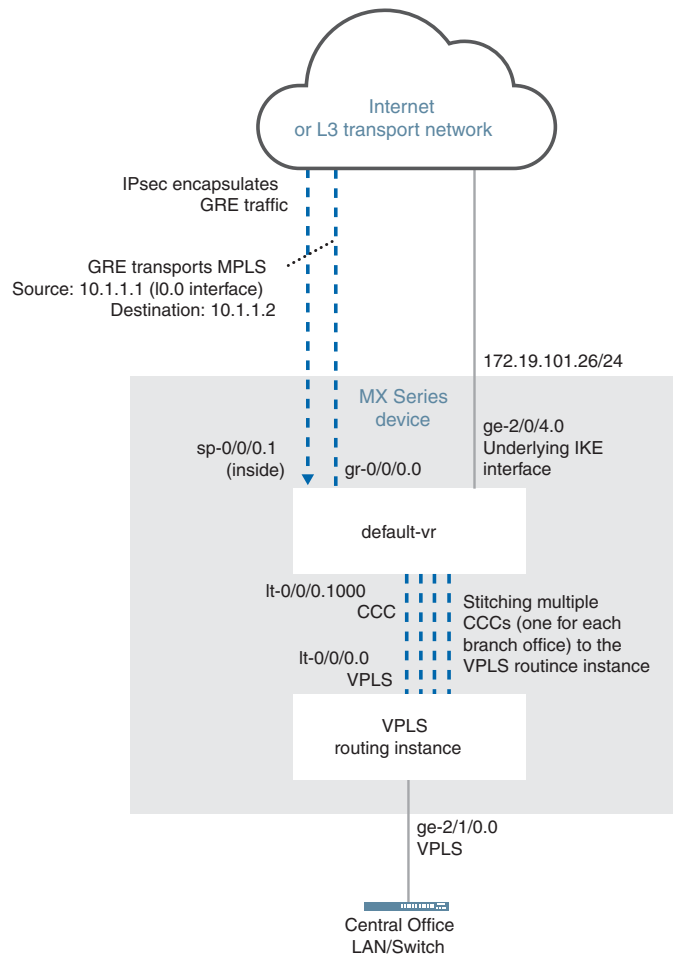
The pseudowires are stitched to a VPLS routing instance using logical tunnel (lt) interfaces. It is possible to use an lt interface unit to terminate a CCC connection and connect this unit to a different unit that is part of a VPLS routing instance. The overall result is as if the pseudowires were terminated directly in the VPLS routing instance. [Figure 123 on page 1593](#) illustrates this configuration.

Figure 123: Central Office Ingress (Head-End) Configuration with an SRX Series Device



You can also use an MX Series device as the central office ingress (head-end) to terminate all branch office connections. The differences in the configuration are due to the way IPsec is configured and the fact that on MX Series devices IDP is not required to reassemble the GRE packets; MX Series devices natively support GRE reassembly. With this configuration, you still use lt interfaces to stitch the CCCs between the remote branch offices and the VPLS routing instance as shown in [Figure 124 on page 1594](#).

Figure 124: Central Office Ingress (Head-End) Configuration with an MX Series Device



Configuration

IN THIS SECTION

- [Configuring the SRX240 Device at the Branch Office | 1595](#)
- [Configuring the SRX650 Device at the Central Office | 1601](#)

In this example, we use SRX Series devices and the branch and ingress (head-end) sites will typically be connected to the Internet by Frame-Relay/T1-E1/xDSL/T3/E3 or even Ethernet. A provider MPLS network is not required.

Configuring the SRX240 Device at the Branch Office

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

set interfaces gr-0/0/0 description "GRE tunnel to SRX650"
set interfaces gr-0/0/0 unit 0 clear-dont-fragment-bit
set interfaces gr-0/0/0 unit 0 tunnel source 10.1.1.2
set interfaces gr-0/0/0 unit 0 tunnel destination 10.1.1.1
set interfaces gr-0/0/0 unit 0 tunnel allow-fragmentation
set interfaces gr-0/0/0 unit 0 family inet mtu 2000
set interfaces gr-0/0/0 unit 0 family inet filter input inet-packet-mode
set interfaces gr-0/0/0 unit 0 family mpls mtu 1900
set interfaces gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode
set interfaces lt-0/0/0 unit 0 encapsulation frame-relay
set interfaces lt-0/0/0 unit 0 dlci 16
set interfaces lt-0/0/0 unit 0 peer-unit 1
set interfaces lt-0/0/0 unit 0 family inet
set interfaces lt-0/0/0 unit 0 description "Flow-vr Instance"
set interfaces lt-0/0/0 unit 1 encapsulation frame-relay
set interfaces lt-0/0/0 unit 1 dlci 16
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet filter input inet-packet-mode
set interfaces lt-0/0/0 unit 1 family inet address 10.1.1.2/32
set interfaces ge-0/0/1 encapsulation ethernet-ccc
set interfaces ge-0/0/1 unit 0 description "CCC Interface to customer LAN"
set interfaces ge-0/0/1 unit 0 family ccc filter input ccc-packet-mode
set interfaces ge-0/0/0 unit 0 family inet address 172.19.101.45/24
set interfaces lo0 unit 0 family inet address 10.2.1.2/32
set interfaces st0 unit 0 family inet
set routing-options static route 0.0.0.0/0 next-hop 172.19.101.1
set routing-options static route 10.1.1.1/32 next-hop lt-0/0/0.1
set routing-options static route 10.2.1.1/32 next-hop gr-0/0/0.0
set routing-options router-id 10.2.1.2
set protocols mpls interface gr-0/0/0.0
set protocols ldp interface gr-0/0/0.0
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 10.2.1.1 interface ge-0/0/1.0 virtual-circuit-id 1
set security ike policy SRX650 mode main
set security ike policy SRX650 proposal-set standard
set security ike policy SRX650 pre-shared-key ascii-text "$ABC123"
set security ike gateway SRX650 ike-policy SRX650
set security ike gateway SRX650 address 172.19.101.26

```



```

set security ike gateway SRX650 external-interface ge-0/0/0.0
set security ipsec policy SRX650 proposal-set standard
set security ipsec vpn SRX650 bind-interface st0.0
set security ipsec vpn SRX650 ike gateway SRX650
set security ipsec vpn SRX650 ike ipsec-policy SRX650
set security ipsec vpn SRX650 establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces gr-0/0/0.0
set security zones security-zone untrust interfaces lo0.0
set security zones security-zone untrust interfaces lt-0/0/0.1
set security zones security-zone untrust interfaces ge-0/0/0.0
set security zones security-zone vpn host-inbound-traffic system-services all
set security zones security-zone vpn host-inbound-traffic protocols all
set security zones security-zone vpn interfaces st0.0
set security zones security-zone trust-flow host-inbound-traffic system-services all
set security zones security-zone trust-flow host-inbound-traffic protocols all
set security zones security-zone trust-flow interfaces lt-0/0/0.0
set security policies from-zone trust-flow to-zone vpn policy gre match source-address any
set security policies from-zone trust-flow to-zone vpn policy gre match destination-address any
set security policies from-zone trust-flow to-zone vpn policy gre match application junos-gre
set security policies from-zone trust-flow to-zone vpn policy gre then permit application-services idp
set security idp idp-policy gre-reassembly rulebase-ips rule match-all match application junos-gre
set security idp idp-policy gre-reassembly rulebase-ips rule match-all then action ignore-connection
set security idp active-policy gre-reassembly
set firewall family inet filter inet-packet-mode term control-traffic from protocol tcp
set firewall family inet filter inet-packet-mode term control-traffic from port 22
set firewall family inet filter inet-packet-mode term control-traffic from port 80
set firewall family inet filter inet-packet-mode term control-traffic from port 8080
set firewall family inet filter inet-packet-mode term control-traffic then accept
set firewall family inet filter inet-packet-mode term packet-mode then packet-mode
set firewall family inet filter inet-packet-mode term packet-mode then accept
set firewall family mpls filter mpls-packet-mode term packet-mode then packet-mode
set firewall family mpls filter mpls-packet-mode term packet-mode then accept
set firewall family ccc filter ccc-packet-mode term all then packet-mode
set firewall family ccc filter ccc-packet-mode term all then accept
set routing-instances flow-vr instance-type virtual-router
set routing-instances flow-vr interface lt-0/0/0.0
set routing-instances flow-vr interface st0.0
set routing-instances flow-vr routing-options static route 10.1.1.1/32 next-hop st0.0
set routing-instances flow-vr routing-options static route 10.1.1.2/32 next-hop lt-0/0/0.0
set security flow tcp-session no-syn-check
set security flow tcp-session no-sequence-check

```



```

set groups test security policies from-zone trust-flow to-zone vpn policy all then permit tcp-options
  syn-check-required
set groups test security policies from-zone trust-flow to-zone vpn policy all then permit tcp-options
  sequence-check-required
set security policies apply-groups test

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the SRX240 at the branch office:

1. Configure a GRE tunnel to the central office.

```

[edit interfaces]
user@host# set gr-0/0/0 description "GRE tunnel to SRX650"
user@host# set gr-0/0/0 unit 0 clear-dont-fragment-bit
user@host# set gr-0/0/0 unit 0 tunnel source 10.1.1.2
user@host# set gr-0/0/0 unit 0 tunnel destination 10.1.1.1
user@host# set gr-0/0/0 unit 0 tunnel allow-fragmentation
user@host# set gr-0/0/0 unit 0 family inet mtu 2000
user@host# set gr-0/0/0 unit 0 family inet filter input inet-packet-mode
user@host# set gr-0/0/0 unit 0 family mpls mtu 1900
user@host# set gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode

```

2. Create a logical interface that connects to the default routing instance.

```

[edit interfaces]
user@host# set lt-0/0/0 unit 0 encapsulation frame-relay
user@host# set lt-0/0/0 unit 0 dlci 16
user@host# set lt-0/0/0 unit 0 peer-unit 1
user@host# set lt-0/0/0 unit 0 family inet
user@host# set lt-0/0/0 unit 0 description "Flow-vr Instance"

```

3. Connect the logical tunnel interface to the flow mode virtual router.

```

[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation frame-relay
user@host# set lt-0/0/0 unit 1 dlci 16
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet filter input inet-packet-mode
user@host# set lt-0/0/0 unit 1 family inet address 10.1.1.2/32

```


4. Connect the CCC interface to the branch LAN.

```
[edit interfaces]
user@host# set ge-0/0/1 encapsulation ethernet-ccc
user@host# set ge-0/0/1 unit 0 description "CCC Interface to customer LAN"
user@host# set ge-0/0/1 unit 0 family ccc filter input ccc-packet-mode
```

5. Configure the interface bound to the default virtual router.

```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 family inet address 172.19.101.45/24
```

6. Set the loopback interface to terminate the CCC connection.

```
[edit interfaces]
user@host# set lo0 unit 0 family inet address 10.2.1.2/32
```

7. Bind the IPsec tunnel interface to the flow-mode virtual router.

```
[edit interfaces]
user@host# set st0 unit 0 family inet
```

8. Set a static route address, which will be the default gateway to the Internet.

```
[edit routing-options]
user@host# set static route 0.0.0.0/0 next-hop 172.19.101.1
```

9. Set a static route for the remote GRE tunnel endpoint.

```
[edit routing-options]
user@host# set static route 10.1.1.1/32 next-hop lt-0/0/0.1
```

10. Set a static route for the loopback interface of the SRX650 ingress (head-end) device.

```
[edit routing-options]
user@host# set static route 10.2.1.1/32 next-hop gr-0/0/0.0
```

11. Configure MPLS and the CCC using LDP as the label protocol.


```
[edit]
user@host# set routing-options router-id 10.2.1.2
user@host# set protocols mpls interface gr-0/0/0.0
user@host# set protocols ldp interface gr-0/0/0.0
user@host# set protocols ldp interface lo0.0
user@host# set protocols l2circuit neighbor 10.2.1.1 interface ge-0/0/1.0 virtual-circuit-id 1
```

12. Configure the IPsec tunnel.

NOTE: The underlying IKE interface is not in the same routing instance as the tunnel interface.

```
[edit security]
user@host# set ike policy SRX650 mode main
user@host# set ike policy SRX650 proposal-set standard
user@host# set ike policy SRX650 pre-shared-key ascii-text "$ABC123"
user@host# set ike gateway SRX650 ike-policy SRX650
user@host# set ike gateway SRX650 address 172.19.101.26
user@host# set ike gateway SRX650 external-interface ge-0/0/0.0
user@host# set ipsec policy SRX650 proposal-set standard
user@host# set ipsec vpn SRX650 bind-interface st0.0
user@host# set ipsec vpn SRX650 ike gateway SRX650
user@host# set ipsec vpn SRX650 ike ipsec-policy SRX650
user@host# set ipsec vpn SRX650 establish-tunnels immediately
```

13. Configure security zones.

NOTE: In a production environment, host-inbound traffic should be restricted to only allow the necessary protocols and services.

```
[edit security]
user@host# set zones security-zone untrust host-inbound-traffic system-services all
user@host# set zones security-zone untrust host-inbound-traffic protocols all
user@host# set zones security-zone untrust interfaces gr-0/0/0.0
user@host# set zones security-zone untrust interfaces lo0.0
user@host# set zones security-zone untrust interfaces lt-0/0/0.1
user@host# set zones security-zone untrust interfaces ge-0/0/0.0
```



```

user@host# set zones security-zone vpn host-inbound-traffic system-services all
user@host# set zones security-zone vpn host-inbound-traffic protocols all
user@host# set zones security-zone vpn interfaces st0.0
user@host# set zones security-zone trust-flow host-inbound-traffic system-services all
user@host# set zones security-zone trust-flow host-inbound-traffic protocols all
user@host# set zones security-zone trust-flow interfaces lt-0/0/0.0

```

14. Configure IDP.

```

[edit security]
user@host# set policies from-zone trust-flow to-zone vpn policy gre match source-address any
user@host# set policies from-zone trust-flow to-zone vpn policy gre match destination-address any
user@host# set policies from-zone trust-flow to-zone vpn policy gre match application junos-gre
user@host# set policies from-zone trust-flow to-zone vpn policy gre then permit application-services idp
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-all match application junos-gre
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-all then action ignore-connection
user@host# set idp active-policy gre-reassembly

```

15. Configure packet-mode filters.

```

[edit firewall]
user@host# set family inet filter inet-packet-mode term control-traffic from protocol tcp
user@host# set family inet filter inet-packet-mode term control-traffic from port 22
user@host# set family inet filter inet-packet-mode term control-traffic from port 80
user@host# set family inet filter inet-packet-mode term control-traffic from port 8080
user@host# set family inet filter inet-packet-mode term control-traffic then accept
user@host# set family inet filter inet-packet-mode term packet-mode then packet-mode
user@host# set family inet filter inet-packet-mode term packet-mode then accept
user@host# set family mpls filter mpls-packet-mode term packet-mode then packet-mode
user@host# set family mpls filter mpls-packet-mode term packet-mode then accept
user@host# set family ccc filter ccc-packet-mode term all then packet-mode
user@host# set family ccc filter ccc-packet-mode term all then accept

```

16. Configure the flow-mode virtual router.

```

[edit routing-instances]
user@host# set flow-vr instance-type virtual-router
user@host# set flow-vr interface lt-0/0/0.0
user@host# set flow-vr interface st0.0
user@host# set flow-vr routing-options static route 10.1.1.1/32 next-hop st0.0
user@host# set flow-vr routing-options static route 10.1.1.2/32 next-hop lt-0/0/0.0

```


17. Disable syn check and sequence check to bypass LDP session from syn check and sequence check.

```
[edit ]
user@host# set security flow tcp-session no-syn-check
user@host# set security flow tcp-session no-sequence-check
```

18. Enable syn check and sequence check at the policy level.

```
[edit ]
user@host# set groups test security policies from-zone trust-flow to-zone vpn policy all then permit
    tcp-options syn-check-required
user@host# set groups test security policies from-zone trust-flow to-zone vpn policy all then permit
    tcp-options sequence-check-required
user@host# set security policies apply-groups test
```

Results

From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the SRX650 Device at the Central Office

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/0/0 unit 0 family inet address 172.19.101.26/24
set interfaces gr-0/0/0 unit 0 clear-dont-fragment-bit
set interfaces gr-0/0/0 unit 0 tunnel source 10.1.1.1
set interfaces gr-0/0/0 unit 0 tunnel destination 10.1.1.2
set interfaces gr-0/0/0 unit 0 tunnel allow-fragmentation
set interfaces gr-0/0/0 unit 0 family inet mtu 1500
set interfaces gr-0/0/0 unit 0 family inet filter input inet-packet-mode
set interfaces gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode
set interfaces lt-0/0/0 unit 0 description "VPLS hub port - Interconnect for CCC to SRX240"
set interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set interfaces lt-0/0/0 unit 0 peer-unit 1000
set interfaces lt-0/0/0 unit 1000 description "Stitch to VPLS for CCC to SRX240"
set interfaces lt-0/0/0 unit 1000 encapsulation ethernet-ccc
set interfaces lt-0/0/0 unit 1000 peer-unit 0
set interfaces lt-0/0/0 unit 1000 family ccc filter input ccc-packet-mode
```



```

set interfaces lt-0/0/0 unit 2000 encapsulation frame-relay
set interfaces lt-0/0/0 unit 2000 dlci 1
set interfaces lt-0/0/0 unit 2000 peer-unit 2001
set interfaces lt-0/0/0 unit 2000 family inet
set interfaces lt-0/0/0 unit 2001 encapsulation frame-relay
set interfaces lt-0/0/0 unit 2001 dlci 1
set interfaces lt-0/0/0 unit 2001 peer-unit 2000
set interfaces lt-0/0/0 unit 2001 family inet filter input inet-packet-mode
set interfaces lt-0/0/0 unit 2001 family inet address 10.1.1.1/32
set interfaces ge-0/0/1 unit 0
set interfaces ge-0/0/1 encapsulation ethernet-vpls
set interfaces lo0 unit 0 family inet address 10.2.1.1/32
set interfaces st0 unit 0 family inet
set routing-options static route 10.1.1.2/32 next-hop lt-0/0/0.2001
set routing-options static route 10.2.1.2/32 next-hop gr-0/0/0.0
set protocols mpls interface gr-0/0/0.0
set protocols ldp interface gr-0/0/0.0
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 10.2.1.2 interface lt-0/0/0.1000 virtual-circuit-id 1
set security ike policy SRX mode main
set security ike policy SRX proposal-set standard
set security ike policy SRX pre-shared-key ascii-text "$ABC123"
set security ike gateway SRX240-1 ike-policy SRX
set security ike gateway SRX240-1 address 172.19.101.45
set security ike gateway SRX240-1 external-interface ge-0/0/0.0
set security ipsec policy SRX proposal-set standard
set security ipsec vpn SRX240-1 bind-interface st0.0
set security ipsec vpn SRX240-1 ike gateway SRX240-1
set security ipsec vpn SRX240-1 ike ipsec-policy SRX
set security ipsec vpn SRX240-1 establish-tunnels immediately
set security zones security-zone untrust host-inbound-traffic system-services all
set security zones security-zone untrust host-inbound-traffic protocols all
set security zones security-zone untrust interfaces lo0.0
set security zones security-zone untrust interfaces lt-0/0/0.2001
set security zones security-zone untrust interfaces gr-0/0/0.0
set security zones security-zone untrust interfaces ge-0/0/0.0
set security zones security-zone vpn host-inbound-traffic system-services all
set security zones security-zone vpn host-inbound-traffic protocols all
set security zones security-zone vpn interfaces st0.0
set security zones security-zone trust-flow host-inbound-traffic system-services all
set security zones security-zone trust-flow host-inbound-traffic protocols all
set security zones security-zone trust-flow interfaces lt-0/0/0.2000
set security policies from-zone trust-flow to-zone vpn policy gre match source-address any
set security policies from-zone trust-flow to-zone vpn policy gre match destination-address any

```



```

set security policies from-zone trust-flow to-zone vpn policy gre match application junos-gre
set security policies from-zone trust-flow to-zone vpn policy gre then permit application-services idp
set security policies from-zone vpn to-zone trust-flow policy gre match source-address any
set security policies from-zone vpn to-zone trust-flow policy gre match destination-address any
set security policies from-zone vpn to-zone trust-flow policy gre match application junos-gre
set security policies from-zone vpn to-zone trust-flow policy gre then permit application-services idp
set security idp idp-policy gre-reassembly rulebase-ips rule match-gre match application junos-gre
set security idp idp-policy gre-reassembly rulebase-ips rule match-gre then action ignore-connection
set security idp active-policy gre-reassembly
set firewall family inet filter inet-packet-mode term control-traffic from protocol tcp
set firewall family inet filter inet-packet-mode term control-traffic from port 22
set firewall family inet filter inet-packet-mode term control-traffic from port 80
set firewall family inet filter inet-packet-mode term control-traffic from port 8080
set firewall family inet filter inet-packet-mode term control-traffic then accept
set firewall family inet filter inet-packet-mode term packet-mode then packet-mode
set firewall family inet filter inet-packet-mode term packet-mode then accept
set firewall family mpls filter mpls-packet-mode term packet-mode then packet-mode
set firewall family mpls filter mpls-packet-mode term packet-mode then accept
set firewall family ccc filter ccc-packet-mode term all then packet-mode
set firewall family ccc filter ccc-packet-mode term all then accept
set routing-instances flow-vr instance-type virtual-router
set routing-instances flow-vr interface lt-0/0/0.2000
set routing-instances flow-vr interface st0.0
set routing-instances flow-vr routing-options static route 10.1.1.1/32 next-hop lt-0/0/0.2000
set routing-instances flow-vr routing-options static route 10.1.1.2/32 next-hop st0.0
set routing-instances vpls-hub instance-type vpls
set routing-instances vpls-hub interface lt-0/0/0.0
set routing-instances vpls-hub interface ge-0/0/1.0
set security flow tcp-session no-syn-check
set security flow tcp-session no-sequence-check
set groups test security policies from-zone trust-flow to-zone vpn policy all then permit tcp-options
    syn-check-required
set groups test security policies from-zone trust-flow to-zone vpn policy all then permit tcp-options
    sequence-check-required
set security policies apply-groups test

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the ingress (head-end) SRX650 device at the central office:

1. Configure the interface bound to the default virtual router.


```
[edit interfaces]
user@host# set ge-0/0/0 unit 0 family inet address 172.19.101.26/24
```

2. Create the GRE tunnel from the SRX650 to the SRX240 device.

NOTE: As the network expands to include multiple branch offices, you will need to add a similar GRE tunnel configuration on the SRX650 device (head-end) along with a corresponding IPsec configuration to connect to each additional SRX device (SRX240).

```
[edit interfaces]
user@host# set gr-0/0/0 unit 0 clear-dont-fragment-bit
user@host# set gr-0/0/0 unit 0 tunnel source 10.1.1.1
user@host# set gr-0/0/0 unit 0 tunnel destination 10.1.1.2
user@host# set gr-0/0/0 unit 0 tunnel allow-fragmentation
user@host# set gr-0/0/0 unit 0 family inet mtu 1500
user@host# set gr-0/0/0 unit 0 family inet filter input inet-packet-mode
user@host# set gr-0/0/0 unit 0 family mpls filter input mpls-packet-mode
```

3. Configure a logical tunnel interface to stitch the CCC connection to the VPLS instance.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 0 description "VPLS hub port - Interconnect for CCC to SRX240"
user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 0 peer-unit 1000
```

4. Set unit 1000 to terminate the CCC connection.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 1000 description "Stitch to VPLS for CCC to SRX240"
user@host# set lt-0/0/0 unit 1000 encapsulation ethernet-ccc
user@host# set lt-0/0/0 unit 1000 peer-unit 0
user@host# set lt-0/0/0 unit 1000 family ccc filter input ccc-packet-mode
```

5. Configure the logical tunnel interface.

```
[edit interfaces]
user@host# set lt-0/0/0 unit 2000 encapsulation frame-relay
```



```

user@host# set lt-0/0/0 unit 2000 dlci 1
user@host# set lt-0/0/0 unit 2000 peer-unit 2001
user@host# set lt-0/0/0 unit 2000 family inet

```

6. Bind the logical tunnel interface to the default virtual router.

```

[edit interfaces]
user@host# set lt-0/0/0 unit 2001 encapsulation frame-relay
user@host# set lt-0/0/0 unit 2001 dlci 1
user@host# set lt-0/0/0 unit 2001 peer-unit 2000
user@host# set lt-0/0/0 unit 2001 family inet filter input inet-packet-mode
user@host# set lt-0/0/0 unit 2001 family inet address 10.1.1.1/32

```

7. Set the interface to the central office LAN network.

```

[edit interfaces]
user@host# set ge-0/0/1 unit 0
user@host# set ge-0/0/1 encapsulation ethernet-vpls

```

8. Set the loopback interface to terminate the CCC connections to each SRX device.

```

[edit interfaces]
user@host# set lo0 unit 0 family inet address 10.2.1.1/32

```

9. Bind the IPsec interface to the flow-mode virtual router.

```

[edit interfaces]
user@host# set st0 unit 0 family inet

```

10. Set a static route for the remote GRE tunnel endpoint.

```

[edit routing-options]
user@host# set static route 10.1.1.2/32 next-hop lt-0/0/0.2001

```

11. Set a static route for the loopback interface of the SRX device.

```

[edit]
user@host# set routing-options static route 10.2.1.2/32 next-hop gr-0/0/0.0

```


12. Configure MPLS and CCC using LDP as the label protocol.

```
[edit protocols]
user@host# set mpls interface gr-0/0/0.0
user@host# set ldp interface gr-0/0/0.0
user@host# set ldp interface lo0.0
user@host# set l2circuit neighbor 10.2.1.2 interface lt-0/0/0.1000 virtual-circuit-id 1
```

13. Configure the IPsec tunnel.

NOTE: The underlying IKE interface is not in the same routing instance as the tunnel interface.

```
[edit security]
user@host# set ike policy SRX mode main
user@host# set ike policy SRX proposal-set standard
user@host# set ike policy SRX pre-shared-key ascii-text "$ABC123"
user@host# set ike gateway SRX240-1 ike-policy SRX
user@host# set ike gateway SRX240-1 address 172.19.101.45
user@host# set ike gateway SRX240-1 external-interface ge-0/0/0.0
user@host# set ipsec policy SRX proposal-set standard
user@host# set ipsec vpn SRX240-1 bind-interface st0.0
user@host# set ipsec vpn SRX240-1 ike gateway SRX240-1
user@host# set ipsec vpn SRX240-1 ike ipsec-policy SRX
user@host# set ipsec vpn SRX240-1 establish-tunnels immediately
```

14. Configure security zones.

NOTE: In a production environment, restrict host-inbound traffic to only the necessary protocols and services.

```
[edit security]
user@host# set zones security-zone untrust host-inbound-traffic system-services all
user@host# set zones security-zone untrust host-inbound-traffic protocols all
user@host# set zones security-zone untrust interfaces lo0.0
user@host# set zones security-zone untrust interfaces lt-0/0/0.2001
user@host# set zones security-zone untrust interfaces gr-0/0/0.0
```



```

user@host# set zones security-zone untrust interfaces ge-0/0/0.0
user@host# set zones security-zone vpn host-inbound-traffic system-services all
user@host# set zones security-zone vpn host-inbound-traffic protocols all
user@host# set zones security-zone vpn interfaces st0.0
user@host# set zones security-zone trust-flow host-inbound-traffic system-services all
user@host# set zones security-zone trust-flow host-inbound-traffic protocols all
user@host# set zones security-zone trust-flow interfaces lt-0/0/0.2000

```

15. Configure IDP.

```

[edit security]
user@host# set policies from-zone trust-flow to-zone vpn policy GRE match source-address any
user@host# set policies from-zone trust-flow to-zone vpn policy GRE match destination-address any
user@host# set policies from-zone trust-flow to-zone vpn policy GRE match application junos-gre
user@host# set policies from-zone trust-flow to-zone vpn policy GRE then permit application-services idp
user@host# set policies from-zone vpn to-zone trust-flow policy GRE match source-address any
user@host# set policies from-zone vpn to-zone trust-flow policy GRE match destination-address any
user@host# set policies from-zone vpn to-zone trust-flow policy GRE match application junos-gre
user@host# set policies from-zone vpn to-zone trust-flow policy GRE then permit application-services idp
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-gre match application junos-gre
user@host# set idp idp-policy gre-reassembly rulebase-ips rule match-gre then action ignore-connection
user@host# set idp active-policy gre-reassembly

```

16. Configure packet-mode filters.

```

[edit firewall]
user@host# set family inet filter inet-packet-mode term control-traffic from protocol tcp
user@host# set family inet filter inet-packet-mode term control-traffic from port 22
user@host# set family inet filter inet-packet-mode term control-traffic from port 80
user@host# set family inet filter inet-packet-mode term control-traffic from port 8080
user@host# set family inet filter inet-packet-mode term control-traffic then accept
user@host# set family inet filter inet-packet-mode term packet-mode then packet-mode
user@host# set family inet filter inet-packet-mode term packet-mode then accept
user@host# set family mpls filter mpls-packet-mode term packet-mode then packet-mode
user@host# set family mpls filter mpls-packet-mode term packet-mode then accept
user@host# set family ccc filter ccc-packet-mode term all then packet-mode
user@host# set family ccc filter ccc-packet-mode term all then accept

```

17. Configure the flow-mode virtual router.

```

[edit routing-instances]

```



```

user@host# set flow-vr instance-type virtual-router
user@host# set flow-vr interface lt-0/0/0.2000
user@host# set flow-vr interface st0.0
user@host# set flow-vr routing-options static route 10.1.1.1/32 next-hop lt-0/0/0.2000
user@host# set flow-vr routing-options static route 10.1.1.2/32 next-hop st0.0

```

18. Configure the VPLS instance.

```

[edit routing-instances]
user@host# set vpls-hub instance-type vpls
user@host# set vpls-hub interface lt-0/0/0.0
user@host# set vpls-hub interface ge-0/0/1.0

```

19. Disable syn check and sequence check to bypass LDP session from syn check and sequence check.

```

[edit ]
user@host# set security flow tcp-session no-syn-check
user@host# set security flow tcp-session no-sequence-check

```

20. Enable syn check and sequence check at the policy level.

```

[edit ]
user@host# set groups test security policies from-zone trust-flow to-zone vpn policy all then permit
    tcp-options syn-check-required
user@host# set groups test security policies from-zone trust-flow to-zone vpn policy all then permit
    tcp-options sequence-check-required
user@host# set security policies apply-groups test

```

Results

From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying Interfaces | 1609](#)
- [Verifying an IPsec tunnel | 1609](#)
- [Verifying GRE | 1609](#)
- [Verifying the CCC/L2 circuit. | 1609](#)
- [Verifying that LDP sessions are working. | 1610](#)

Confirm that the configuration is working properly.

Verifying Interfaces

Purpose

Verify that the interfaces are configured properly on each device in the VPLS network.

Action

From configuration mode, enter **show interfaces** and verify that the IP addressing is correct for each interface, including logical tunnel (lt), loopback (lo), GRE (gr), IPsec tunnel st0, and GE interfaces.

Verifying an IPsec tunnel

Purpose

Verify that an IPsec tunnel is working.

Action

From operational mode, enter the **show security ipsec security associations** and the **show security ipsec statistics** command.

Verifying GRE

Purpose

Verify that GRE is working.

Action

From operational mode, enter the **show security flow session protocol gre** command. You can also do a ping between loopback addresses.

Verifying the CCC/L2 circuit.

Purpose

Verify that the CCC/L2 circuit is working.

Action

From operational mode, enter the **show connections** command.

Verifying that LDP sessions are working.

Purpose

Verify that LDP sessions are being created between devices.

Action

From operational mode, enter the **show interfaces gr-0/0/0 detail** command.

RELATED DOCUMENTATION

[VPLS Overview | 1546](#)

[Understanding VPLS Interfaces | 1561](#)

Understanding Selective Stateless Packet-Based Services

[MPLS Overview | 8](#)

Example: Configuring VPLS with BGP Signaling

IN THIS SECTION

- [Requirements | 1610](#)
- [Overview | 1611](#)
- [Configuration | 1612](#)
- [Verification | 1625](#)

This example shows how to configure VPLS with BGP signaling between two devices.

Requirements


Before you begin, see *Understanding Selective Stateless Packet-Based Services* .

Overview

This example shows a minimum configuration for PE devices and CE devices to create a VPLS network with BGP signaling. The topology consists of two PE devices and two CE devices. In this example, you configure a VPLS routing instance vpls-instance between two PE devices, PE1 and PE2. You also configure the CE1 and CE2 devices that use Ethernet-based interfaces to connect VLAN 600 to their local PE devices. On the CE1 device, configure the Fast Ethernet interface that connects to the PE1 device. The VLAN identifier and IP address must match those of the CE2 device.

Figure 125 on page 1611 shows the topology used in this example.

Figure 125: Configuring VPLS with BGP Signaling



Configuration

IN THIS SECTION

- [Configuring the CE1 Device | 1612](#)
- [Configuring the PE1 Device | 1613](#)
- [Configuring the PE2 Device | 1618](#)
- [Configuring the CE2 Device | 1624](#)

Configuring the CE1 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces fe-0/0/3 vlan-tagging
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family inet address 10.11.3.1/24
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Enable VLAN tagging on the VPLS interface.

```
[edit interfaces fe-0/0/3]
user@host# set vlan-tagging
```

2. Configure the VLAN ID on the logical interface.

```
[edit interfaces fe-0/0/3 unit 0]
user@host# set vlan-id 600
```

3. Configure the VPLS family on the logical interface.

```
[edit interfaces fe-0/0/3 unit 0]
```



```
user@host# set family inet address 10.11.3.1/24
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
fe-0/0/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 600;
    family inet {
      address 10.11.3.1/24;
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the PE1 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set system host-name PE1
set interfaces fe-0/0/3 description "CE1 on PE1"
set interfaces fe-0/0/3 vlan-tagging
set interfaces fe-0/0/3 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family vpls
set interfaces fe-0/0/5 vlan-tagging
set interfaces fe-0/0/5 unit 37 vlan-id 37
set interfaces fe-0/0/5 unit 37 family inet address 172.28.2.133/30
set interfaces fe-0/0/5 unit 37 family mpls
set interfaces lo0 unit 0 family inet address 172.28.1.2/32
set routing-options router-id 172.28.1.2
set routing-options autonomous-system 65512
set protocols rsvp interface fe-0/0/5.37
set protocols mpls label-switched-path pe1-to-pe2 to 172.28.1.1
```



```

set protocols mpls interface fe-0/0/5.37
set protocols mpls interface lo0.0
set protocols bgp group vpls-peering type internal
set protocols bgp group vpls-peering local-address 172.28.1.2
set protocols bgp group vpls-peering family l2vpn signaling
set protocols bgp group vpls-peering neighbor 172.28.1.1
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/5.37
set routing-instances vpls-instance description "Routing instance from VPLS routing"
set routing-instances vpls-instance instance-type vpls
set routing-instances vpls-instance interface fe-0/0/3.0
set routing-instances vpls-instance route-distinguisher 172.28.1.2:1
set routing-instances vpls-instance vrf-target target:65512:1
set routing-instances vpls-instance protocols vpls site-range 10
set routing-instances vpls-instance protocols vpls no-tunnel-services site site10 automatic-site-id

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure PE1:

1. Configure the hostname for the PE1 device.

```

[edit ]
user@host# set system host-name PE1

```

2. Configure VPLS VLAN encapsulation on the VPLS PE1 device.

```

[edit interfaces]
user@host# set fe-0/0/3 description "CE1 on PE1"
user@host# set fe-0/0/3 vlan-tagging
user@host# set fe-0/0/3 encapsulation vlan-vpls
user@host# set fe-0/0/3 unit 0 encapsulation vlan-vpls
user@host# set fe-0/0/3 unit 0 vlan-id 600
user@host# set fe-0/0/3 unit 0 family vpls

```

3. Configure the routing interface on the VPLS PE1 device.

```

[edit interfaces]
user@host# set fe-0/0/5 vlan-tagging
user@host# set fe-0/0/5 unit 37 vlan-id 37

```



```

user@host# set fe-0/0/5 unit 37 family inet address 172.28.2.133/30
user@host# set fe-0/0/5 unit 37 family mpls
user@host# set lo0 unit 0 family inet address 172.28.1.2/32

```

NOTE: For this example, it is optional to configure VLAN tagging. Remove the VLAN tagging configuration on the physical interfaces if you do not plan to configure VLAN tagging.

4. Configure the routing options on the VPLS PE1 device.

```

[edit routing-options]
user@host# set router-id 172.28.1.2
user@host# set autonomous-system 65512

```

5. Configure RSVP on the VPLS PE1 device.

```

[edit protocols]
user@host# set rsvp interface fe-0/0/5.37

```

6. Configure MPLS on the VPLS PE1 device.

```

[edit protocols]
user@host# set mpls label-switched-path pe1-to-pe2 to 172.28.1.1
user@host# set mpls interface fe-0/0/5.37
user@host# set mpls interface lo0.0

```

7. Configure BGP on the VPLS PE1 device.

```

[edit protocols]
user@host# set bgp group vpls-peering type internal
user@host# set bgp group vpls-peering local-address 172.28.1.2
user@host# set bgp group vpls-peering family l2vpn signaling
user@host# set bgp group vpls-peering neighbor 172.28.1.1

```

8. (Optional) Configure OSPF on the VPLS PE1 device.

NOTE: For this example, it is optional to configure OFPF. You must configure OSPF only in cases where two PE devices are not connected directly.

```
[edit protocols]
user@host# set ospf area 0.0.0.0 interface lo0.0 passive
user@host# set ospf area 0.0.0.0 interface fe-0/0/5.37
```

9. Create a VPLS routing instance.

```
[edit ]
user@host# set routing-instances vpls-instance
```

10. Configure a VPLS routing instance.

```
[edit routing-instances vpls-instance]
user@host# set description "Routing instance from VPLS routing"
user@host# set instance-type vpls
user@host# set interface fe-0/0/3.0
user@host# set route-distinguisher 172.28.1.2:1
user@host# set vrf-target target:65512:1
user@host# set protocols vpls site-range 10
user@host# set protocols vpls no-tunnel-services site site10 automatic-site-id
```

Results

From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show system
host-name PE1;
```

```
[edit]
user@host# show interfaces
fe-0/0/5 {
  vlan-tagging;
  unit 37 {
    vlan-id 37;
```



```

        family inet {
            address 172.28.2.133/30;
        }
        family mpls;
    }
}
fe-0/0/3 {
    description "CE1 on PE1";
    vlan-tagging;
    encapsulation vlan-vpls;
    unit 0 {
        encapsulation vlan-vpls;
        vlan-id 600;
        family vpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.28.1.2/32;
        }
    }
}
}

```

```

[edit]
user@host# show routing-options
router-id 172.28.1.2;
autonomous-system 65512;

```

```

[edit]
user@host# show protocols
rsvp {
    interface fe-0/0/5.37;
}
mpls {
    label-switched-path pe1-to-pe2 {
        to 172.28.1.1;
    }
    interface fe-0/0/5.37;
    interface lo0.0;
}
bgp {
    group vpls-peering {

```



```

    type internal;
    local-address 172.28.1.2;
    family l2vpn {
        signaling;
    }
    neighbor 172.28.1.1;
}
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface fe-0/0/5.37;
    }
}
}

```

```

[edit]
user@host# show routing-instances
vpls-instance {
    description "Routing instance from VPLS routing";
    instance-type vpls;
    interface fe-0/0/3.0;
    route-distinguisher 172.28.1.2:1;
    vrf-target target:65512:1;
    protocols {
        vpls {
            site-range 10;
            no-tunnel-services;
            site site10 {
                automatic-site-id;
            }
        }
    }
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the PE2 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.


```

set system host-name PE2
set interfaces fe-0/0/3 description "CE2 on PE2"
set interfaces fe-0/0/3 vlan-tagging
set interfaces fe-0/0/3 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 encapsulation vlan-vpls
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family vpls
set interfaces fe-0/0/5 vlan-tagging
set interfaces fe-0/0/5 unit 37 vlan-id 37
set interfaces fe-0/0/5 unit 37 family inet address 172.28.2.133/30
set interfaces fe-0/0/5 unit 37 family mpls
set interfaces lo0 unit 0 family inet address 172.28.1.1/32
set routing-options router-id 172.28.1.1
set routing-options autonomous-system 65512
set protocols rsvp interface fe-0/0/5.37
set protocols mpls label-switched-path pe2-to-pe1 to 172.28.1.2
set protocols mpls interface fe-0/0/5.37
set protocols mpls interface lo0.0
set protocols bgp group vpls-peering type internal
set protocols bgp group vpls-peering local-address 172.28.1.1
set protocols bgp group vpls-peering family l2vpn signaling
set protocols bgp group vpls-peering neighbor 172.28.1.2
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-0/0/5.37
set routing-instances vpls-instance description "Routing instance for VPLS routing"
set routing-instances vpls-instance instance-type vpls
set routing-instances vpls-instance interface fe-0/0/3.0
set routing-instances vpls-instance route-distinguisher 172.28.1.1:1
set routing-instances vpls-instance vrf-target target:65512:1
set routing-instances vpls-instance protocols vpls site-range 10
set routing-instances vpls-instance protocols vpls no-tunnel-services site site11 automatic-site-id

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure PE2:

1. Configure the hostname for the device.

```

[edit ]
user@host# set system host-name PE2

```

2. Configure VPLS VLAN encapsulation on the VPLS PE2 device.


```
[edit interfaces]
user@host# set fe-0/0/3 description "CE2 on PE2"
user@host# set fe-0/0/3 vlan-tagging
user@host# set fe-0/0/3 encapsulation vlan-vpls
user@host# set fe-0/0/3 unit 0 encapsulation vlan-vpls
user@host# set fe-0/0/3 unit 0 vlan-id 600
user@host# set fe-0/0/3 unit 0 family vpls
```

3. Configure the routing interface on the VPLS PE2 device.

```
[edit interfaces]
user@host# set fe-0/0/5 vlan-tagging
user@host# set fe-0/0/5 unit 37 vlan-id 37
user@host# set fe-0/0/5 unit 37 family inet address 172.28.2.133/30
user@host# set fe-0/0/5 unit 37 family mpls
user@host# set lo0 unit 0 family inet address 172.28.1.1/32
```

NOTE: For this example, it is optional to configure VLAN tagging. Remove the VLAN tagging configuration on the physical interfaces if you do not plan to configure VLAN tagging.

4. Configure the routing options on the VPLS PE2 device.

```
[edit routing-options]
user@host# set router-id 172.28.1.1
user@host# set autonomous-system 65512
```

5. Configure RSVP on the VPLS PE2 device.

```
[edit protocols]
user@host# set rsvp interface fe-0/0/5.37
```

6. Configure MPLS on the VPLS PE2 device.

```
[edit protocols]
user@host# set mpls label-switched-path pe2-to-pe1 to 172.28.1.2
user@host# set mpls interface fe-0/0/5.37
```



```
user@host# set mpls interface lo0.0
```

7. Configure BGP on the VPLS PE2 device.

```
[edit protocols]
user@host# set bgp group vpls-peering type internal
user@host# set bgp group vpls-peering local-address 172.28.1.1
user@host# set bgp group vpls-peering family l2vpn signaling
user@host# set bgp group vpls-peering neighbor 172.28.1.2
```

8. (Optional) Configure OSPF on the VPLS PE2 device.

NOTE: For this example, it is optional to configure OSPF. You must configure OSPF only in cases where two PE devices are not connected directly.

```
[edit protocols]
user@host# set ospf area 0.0.0.0 interface lo0.0 passive
user@host# set ospf area 0.0.0.0 interface fe-0/0/5.37
```

9. Create a VPLS routing instance.

```
[edit ]
user@host# set routing-instances vpls-instance
```

10. Configure a VPLS routing instance.

```
[edit routing-instances vpls-instance]
user@host# set description "Routing instance for VPLS routing"
user@host# set instance-type vpls
user@host# set interface fe-0/0/3.0
user@host# set route-distinguisher 172.28.1.1:1
user@host# set vrf-target target:65512:1
user@host# set protocols vpls site-range 10
user@host# set protocols vpls no-tunnel-services site site11 automatic-site-id
```

Results

From configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show system
host-name PE2;
```

```
[edit]
user@host# show interfaces
fe-0/0/5 {
  vlan-tagging;
  unit 37 {
    vlan-id 37;
    family inet {
      address 172.28.2.133/30;
    }
    family mpls;
  }
}
fe-0/0/3 {
  description "CE2 on PE2";
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 0 {
    encapsulation vlan-vpls;
    vlan-id 600;
    family vpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 172.28.1.1/32;
    }
  }
}
```

```
[edit]
user@host# show routing-options
router-id 172.28.1.1;
autonomous-system 65512;
```



```

[edit]
user@host# show protocols
rsvp {
    interface fe-0/0/5.37;
}
mpls {
    label-switched-path pe2-to-pe1 {
        to 172.28.1.2;
    }
    interface fe-0/0/5.37;
    interface lo0.0;
}
bgp {
    group vpls-peering {
        type internal;
        local-address 172.28.1.1;
        family l2vpn {
            signaling;
        }
        neighbor 172.28.1.1;
    }
}
ospf {
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface fe-0/0/5.37;
    }
}

```

```

[edit]
user@host# show routing-instances
vpls-instance {
    description "Routing instance from VPLS routing";
    instance-type vpls;
    interface fe-0/0/3.0;
    route-distinguisher 172.28.1.1:1;
    vrf-target target:65512:1;
    protocols {
        vpls {
            site-range 10;
            no-tunnel-services;
            site site11 {

```



```

        automatic-site-id;
    }
}
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring the CE2 Device

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

set interfaces fe-0/0/3 vlan-tagging
set interfaces fe-0/0/3 unit 0 vlan-id 600
set interfaces fe-0/0/3 unit 0 family inet address 10.11.3.2/24

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Enable VLAN tagging on the VPLS interface.

```

[edit interfaces fe-0/0/3]
user@host# set vlan-tagging

```

2. Configure the VLAN ID on the logical interface.

```

[edit interfaces fe-0/0/3 unit 0]
user@host# set vlan-id 600

```

3. Configure the VPLS family on the logical interface.

```

[edit interfaces fe-0/0/3 unit 0]
user@host# set family inet address 10.11.3.2/24

```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
fe-0/0/3 {  
  vlan-tagging;  
  unit 0 {  
    vlan-id 600;  
    family inet {  
      address 10.11.3.2/24;  
    }  
  }  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

NOTE: If VLAN trunking is not needed between the CE devices, remove the configuration on VLAN tagging on the interfaces connecting the CE and PE devices. Also, use ethernet-VPLS-encapsulation instead of vlan-vpls on the CE facing interfaces of the PE devices.

Verification

IN THIS SECTION

- [Verifying Interfaces | 1625](#)
- [Verifying Routing Information | 1626](#)
- [Verifying VPLS Information | 1626](#)
- [Verifying Automatic Site Identifier Generation | 1626](#)

Confirm that the configuration is working properly.

Verifying Interfaces

Purpose

Verify that the interfaces are configured correctly.

Action

From operational mode, enter the **show interfaces terse** command.

Verifying Routing Information

Purpose

Verify that the routing information is configured correctly.

Action

From operational mode, enter the following commands:

- **show route forwarding-table family mpls**
- **show route forwarding-table family vpls (destination | extensive | matching | table)**
- **show route instance (detail)**

Verifying VPLS Information

Purpose

Verify that the VPLS is configured correctly.

Action

From operational mode, enter the following commands:

- **show system statistics vpls**
- **show vpls connections**
- **show vpls statistics**

Verifying Automatic Site Identifier Generation

Purpose

Verify that the automatic site identifier has been generated.

Action

From operational mode, enter the **show vpls connections** command.

```
[edit]
user@host# show vpls connections
Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down NP -- interface hardware not present
CM -- control-word mismatch -> -- only outbound connection is up
CN -- circuit not provisioned <- -- only inbound connection is up
OR -- out of range Up -- operational
```


OL -- no outgoing label Dn -- down
 LD -- local site signaled down CF -- call admission control failure
 RD -- remote site signaled down SC -- local and remote site ID collision
 LN -- local site not designated LM -- local site ID not minimum designated
 RN -- remote site not designated RM -- remote site ID not minimum designated
 XX -- unknown connection status IL -- no incoming label
 MM -- MTU mismatch MI -- Mesh-Group ID not available
 BK -- Backup connection ST -- Standby connection
 PF -- Profile parse failure PB -- Profile busy
 RS -- remote site standby SN -- Static Neighbor
 VM -- VLAN ID mismatch
 Legend for interface status
 Up -- operational
 Dn -- down
 Instance: customer2
 Local site: airwalk (2)
 connection-site Type St Time last up # Up trans
 4 rmt Up Mar 1 03:26:21 2012 1
 Remote PE: 200.100.100.2, Negotiated control-word: No
 Incoming label: 262148, Outgoing label: 262146
 Local interface: lsi.1048838, Status: Up, Encapsulation: VPLS
 Description: Intf - vpls customer2 local site 2 remote site 4
 Instance: customer4
 Local site: airwalk (6)
 connection-site Type St Time last up # Up trans
 8 rmt Up Feb 21 03:27:33 2012 1
 Remote PE: 200.200.200.2, Negotiated control-word: No
 Incoming label: 262160, Outgoing label: 262174
 Local interface: lsi.1048836, Status: Up, Encapsulation: VPLS
 Description: Intf - vpls customer4 local site 6 remote site 8

RELATED DOCUMENTATION

[VPLS Overview](#) | 1546

[Understanding VPLS Interfaces](#) | 1561

[MPLS Overview](#) | 8

Example: Configuring BGP on the VPLS PE Router

IN THIS SECTION

- [Requirements | 1628](#)
- [Overview | 1628](#)
- [Configuration | 1629](#)
- [Verification | 1629](#)

This example shows how to configure BGP on the VPLS PE router.

Requirements

Before you begin:

- See *Understanding Selective Stateless Packet-Based Services* .
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE. See [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#). Then configure MPLS LSPs on the PE routers. See [“Example: Configuring MPLS on the VPLS PE Router” on page 1585](#). Alternatively, configure LDP on the PE routers. See [“Example: Configuring LDP on the VPLS PE Router” on page 1587](#).
- Configure routing options on the PE router. See [“Example: Configuring Routing Options on the VPLS PE Router” on page 1630](#).

Overview

In this example, you configure an internal BGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. The PE routers use this information to determine which labels to use for traffic destined for remote sites.

NOTE: On all SRX Series devices, BGP-based virtual private LAN service (VPLS) works on child ports and physical interfaces, but not over aggregated Ethernet (ae) interfaces.

Configuration

Step-by-Step Procedure

To configure BGP on the VPLS PE router:

1. Configure the BGP internal group on the VPLS PE router.

```
[edit ]
user@host# set protocols bgp group ibgp type internal local-address 10.255.7.168 neighbor 10.255.7.164
```

2. Configure the BGP family L2vpn and specify NLRI signaling.

```
[edit ]
user@host# set protocols bgp family L2 VPN signaling
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show protocols** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring Routing Options on the VPLS PE Router

IN THIS SECTION

- Requirements | 1630
- Overview | 1630
- Configuration | 1630
- Verification | 1631

This example shows how to configure the routing options on the VPLS PE router.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE, see [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#) and then MPLS LSPs on the PE routers, see [“Example: Configuring MPLS on the VPLS PE Router” on page 1585](#). Alternatively configure LDP on the PE routers, see [“Example: Configuring LDP on the VPLS PE Router” on page 1587](#).

Overview

This example describes how to specify the router ID and the AS number for each router involved in the VPLS. In this example, the routers PE1 and PE2 use the same AS number (100).

Configuration

Step-by-Step Procedure

To configure the routing options on the VPLS PE router:

1. Configure the router ID on the VPLS PE router.

```
[edit]
user@host# set routing-options router-id 10.255.7.168
```

2. Configure the AS number on the VPLS PE router.

```
[edit]
user@host# set routing-options autonomous-system 100
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show routing-options** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Understanding VPLS VLAN Encapsulation

Gigabit Ethernet IQ, Gigabit Ethernet PIMs with small form-factor pluggable optics (SFPs), SRX Series devices with Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with VLAN tagging enabled can use flexible Ethernet services, VLAN virtual private LAN service (VPLS) encapsulation.

NOTE: VLAN encapsulation is not supported on SRX100 devices because there is no Gigabit Ethernet port.

Aggregated Ethernet interfaces configured for VPLS can use Ethernet VPLS or VLAN VPLS.

To configure the encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface, include the **encapsulation** statement at the **[edit interfaces interface-name]** hierarchy level, specifying **vlan-ccc** or **vlan-vpls**:

```
[edit interfaces interface-name] encapsulation (vlan-ccc | vlan-vpls);
```

To configure the encapsulation on an aggregated Ethernet interface, include the encapsulation statement at the **[edit interfaces interface-name]** hierarchy level, specifying **ethernet-vpls** or **vlan-vpls**:

```
[edit interfaces interface-name] encapsulation (ethernet-vpls | vlan-vpls);
```

Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.

RELATED DOCUMENTATION

[Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces | 1637](#)

[Understanding VPLS VLAN Encapsulation on a Logical Interface | 1632](#)

[Example: Configuring Extended VLAN VPLS Encapsulation | 1639](#)

[Example: Configuring VPLS VLAN Encapsulation | 1633](#)

[VPLS Overview | 1546](#)

Understanding VPLS VLAN Encapsulation on a Logical Interface

You can configure a logical interface with VLAN VPLS encapsulation by using the following methods:

- Configure the physical interface with the same encapsulation and set VLAN ID of 512 or higher.
- Configure the physical interface with flexible Ethernet services encapsulation. If you configure flexible Ethernet services encapsulation, the VLAN ID restriction is removed.

Ethernet interfaces in VLAN mode can have multiple logical interfaces. In VPLS mode, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for VPLS VLAN. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for VPLS VLAN.

For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)[Understanding VPLS VLAN Encapsulation | 1631](#)

Example: Configuring VPLS VLAN Encapsulation

IN THIS SECTION

- [Requirements | 1633](#)
- [Overview | 1634](#)
- [Configuration | 1634](#)
- [Verification | 1636](#)

This example shows how to configure VPLS VLAN encapsulation and enable it on the physical and the logical interfaces.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE, see [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#) and then MPLS LSPs on the PE routers, see [“Example: Configuring MPLS on the VPLS PE Router” on page 1585](#). Alternatively configure LDP on the PE routers, see [“Example: Configuring LDP on the VPLS PE Router” on page 1587](#).

- Configure routing options on the PE router. See [“Example: Configuring Routing Options on the VPLS PE Router” on page 1630](#).
- Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See [“Example: Configuring BGP on the VPLS PE Router” on page 1628](#).

Overview

This example describes how to enable VLAN tagging on VPLS interface ge-3/0/6, configure the encapsulation type on the physical and logical interfaces, and configure the VPLS family on the logical interface.

NOTE: Perform the following CLI quick configuration and procedures on all of the PE interfaces (CE facing).

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-3/0/6 vlan-tagging
set interfaces ge-3/0/6 encapsulation vlan-vpls
set interfaces ge-3/0/6 unit 0 encapsulation vlan-vpls
set interfaces ge-3/0/6 unit 0 vlan-id 512
set interfaces ge-3/0/6 unit 0 family vpls
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure VPLS VLAN encapsulation:

1. Enable VLAN tagging on the VPLS interface.

```
[edit interfaces ge-3/0/6]
user@host# set vlan-tagging
```


2. Configure the encapsulation type on the physical interface.

```
[edit interfaces ge-3/0/6]  
user@host# set interfaces ge-3/0/6 encapsulation vlan-vpls
```

3. Configure the encapsulation type on the logical interface.

```
[edit interfaces ge-3/0/6 unit 0]  
user@host# set encapsulation vlan-vpls
```

4. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-3/0/6 unit 0]  
user@host# set vlan-id 512
```

5. Configure the family VPLS on the logical interface.

```
[edit interfaces ge-3/0/6 unit 0]  
user@host# set family vpls
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces ge-3/0/6** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]  
user@host# show interfaces ge-3/0/6  
vlan-tagging;  
encapsulation vlan-vpls;  
unit 0 {  
    encapsulation vlan-vpls;  
    vlan-id 512;  
    family vpls;  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying VPLS VLAN Encapsulation | 1636](#)
- [Verifying VPLS VLAN Encapsulation for Logical Interfaces | 1636](#)

Confirm that the configuration is working properly.

Verifying VPLS VLAN Encapsulation

Purpose

Verify that the VPLS VLAN encapsulation is enabled at the interfaces.

Action

From operational mode, enter the **show interfaces** command.

Verifying VPLS VLAN Encapsulation for Logical Interfaces

Purpose

Verify that the VPLS VLAN encapsulation is enabled at the logical interface.

Action

From operational mode, enter the **show interfaces ge-3/0/6 unit 0** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring VPLS VLAN Encapsulation on Gigabit Ethernet Interfaces

IN THIS SECTION

- Requirements | 1637
- Overview | 1638
- Configuration | 1638
- Verification | 1638

This example shows how to configure the VPLS VLAN encapsulation on either a Gigabit Ethernet IQ or Gigabit Ethernet physical interface.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE, see [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#) and then MPLS LSPs on the PE routers, see [“Example: Configuring MPLS on the VPLS PE Router” on page 1585](#). Alternatively configure LDP on the PE routers, see [“Example: Configuring LDP on the VPLS PE Router” on page 1587](#).
- Configure routing options on the PE router. See [“Example: Configuring Routing Options on the VPLS PE Router” on page 1630](#).
- Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See [“Example: Configuring BGP on the VPLS PE Router” on page 1628](#).

Overview

This example describes how to configure Ethernet VPLS encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface and enable the VPLS family on the interface.

Configuration

Step-by-Step Procedure

To configure VPLS VLAN encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface:

1. Configure the ethernet-vpls encapsulation on the interface.

```
[edit ]
user@host# set interfaces ge-3/0/6 encapsulation ethernet-vpls
```

2. Enable the VPLS family on the interface.

```
[edit ]
user@host# set interfaces ge-3/0/6 unit 0 family vpls
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify the configuration is working properly, enter the **show interfaces** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Example: Configuring Extended VLAN VPLS Encapsulation

IN THIS SECTION

- [Requirements | 1639](#)
- [Overview | 1640](#)
- [Configuration | 1640](#)
- [Verification | 1641](#)

This example shows how to configure extended VLAN VPLS encapsulation and enable it on the physical and the logical interfaces.

Requirements

Before you begin:

- Before you begin, see *Understanding Selective Stateless Packet-Based Services*.
- Configure the interfaces that will carry the VPLS traffic between the PE router and the CE devices. See [“Example: Configuring Routing Interfaces on the VPLS PE Router” on page 1564](#) and [“Example: Configuring the Interface to the VPLS CE Device” on page 1565](#).
- Create a VPLS routing instance on each PE router that is participating in the VPLS. See [“Example: Configuring the VPLS Routing Instance” on page 1576](#).
- Configure an IGP on the PE routers to exchange routing information. See [“Example: Configuring OSPF on the VPLS PE Router” on page 1581](#).
- Configure RSVP-TE, see [“Example: Configuring RSVP on the VPLS PE Router” on page 1583](#) and then MPLS LSPs on the PE routers, see [“Example: Configuring MPLS on the VPLS PE Router” on page 1585](#). Alternatively configure LDP on the PE routers, see [“Example: Configuring LDP on the VPLS PE Router” on page 1587](#).
- Configure routing options on the PE router. See [“Example: Configuring Routing Options on the VPLS PE Router” on page 1630](#).
- Configure an IBGP session between PE routers so that the routers can exchange information about routes originating and terminating in the VPLS. See [“Example: Configuring BGP on the VPLS PE Router” on page 1628](#).

Overview

This example describes how to enable VLAN tagging on the VPLS interface ge-3/0/6, configure the extended-vlan-vpls type on the physical and logical interfaces, and configure the VPLS family on the logical interface.

NOTE: Perform the following CLI quick configurations and procedures on all PE interfaces (CE facing).

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-3/0/6 vlan-tagging
set interfaces ge-3/0/6 encapsulation extended-vlan-vpls
set interfaces ge-3/0/6 unit 0 vlan-id 100
set interfaces ge-3/0/6 unit 0 family vpls
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure extended VPLS VLAN encapsulation:

1. Enable VLAN tagging on the VPLS interface as it will receive tagged packets from CE.

```
[edit interfaces ge-3/0/6]
user@host# set vlan-tagging
```

2. Configure the encapsulation type on the physical interface.

```
[edit interfaces ge-3/0/6]
user@host# set interfaces ge-3/0/6 encapsulation vlan-vpls
```

3. Configure the VLAN ID on the logical interface.


```
[edit interfaces ge-3/0/6 unit 0]
user@host# set encapsulation vlan-vpls vlan-id 100
```

4. Configure the VPLS family on the logical interface.

```
[edit interfaces ge-3/0/6 unit 0]
user@host# set family vpls
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces ge-3/0/6** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces ge-3/0/6
vlan-tagging;
encapsulation extended-vlan-vpls;
unit 0 {
    encapsulation vlan-vpls;
    vlan-id 100;
    family vpls;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying Extended VLAN VPLS Encapsulation | 1641](#)
- [Verifying Extended VLAN VPLS Encapsulation for Logical Interfaces | 1642](#)

Confirm that the configuration is working properly.

Verifying Extended VLAN VPLS Encapsulation

Purpose

Verify that the extended VLAN VPLS encapsulation is enabled at the interfaces.

Action

From operational mode, enter the **show interfaces** command.

Verifying Extended VLAN VPLS Encapsulation for Logical Interfaces**Purpose**

Verify that the extended VLAN VPLS encapsulation is enabled at the logical interface.

Action

From operational mode, enter the **show interfaces ge-3/0/6 unit 0** command.

RELATED DOCUMENTATION

[VPLS Configuration Overview | 1550](#)

[VPLS Overview | 1546](#)

Configuring Circuit Cross-Connect (CCC) and Translational Cross-Connect (TCC)

IN THIS CHAPTER

- [CCC Overview | 1643](#)
- [Understanding Carrier-of-Carriers VPNs | 1644](#)
- [Understanding Interprovider and Carrier-of-Carriers VPNs | 1647](#)
- [Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\) | 1648](#)
- [VLAN CCC Encapsulation on Transport Side of Pseudowire Client Logical Interfaces Overview | 1651](#)
- [Transmitting Nonstandard BPDUs | 1654](#)
- [TCC Overview | 1654](#)
- [Configuring Layer 2 Switching Cross-Connects Using CCC | 1655](#)
- [Configuring MPLS LSP Tunnel Cross-Connects Using CCC | 1665](#)
- [Configuring TCC | 1671](#)
- [CCC and TCC Graceful Restart | 1677](#)
- [Configuring CCC and TCC Graceful Restart | 1678](#)
- [Configuring an MPLS-Based VLAN CCC Using the Connection Method \(CLI Procedure\) | 1679](#)
- [Configuring CCC Switching for Point-to-Multipoint LSPs | 1681](#)
- [Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\) | 1684](#)
- [Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\) | 1689](#)

CCC Overview

Circuit cross-connect (CCC) allows you to configure transparent connections between two circuits, where a circuit can be a Frame Relay data-link connection identifier (DLCI), an Asynchronous Transfer Mode (ATM) virtual circuit (VC), a Point-to-Point Protocol (PPP) interface, a Cisco High-Level Data Link Control (HDLC) interface, or an MPLS label-switched path (LSP). Using CCC, packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. No other processing—such as header checksums, time-to-live (TTL) decrementing, or protocol processing—is done.

NOTE: The QFX10000 Series switches do not support ATM virtual circuits.

CCC circuits fall into two categories: logical interfaces, which include DLCIs, VCs, virtual local area network (VLAN) IDs, PPP and Cisco HDLC interfaces, and LSPs. The two circuit categories provide three types of cross-connect:

- Layer 2 switching—Cross-connects between logical interfaces provide what is essentially Layer 2 switching. The interfaces that you connect must be of the same type.
- MPLS tunneling—Cross-connects between interfaces and LSPs allow you to connect two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the conduit.
- LSP stitching—Cross-connects between LSPs provide a way to “stitch” together two label-switched paths, including paths that fall in two different traffic engineering database areas.

For Layer 2 switching and MPLS tunneling, the cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first. For LSP stitching, the cross-connect is unidirectional.

You can police (control) the amount of traffic flowing over CCC circuits. For more information, see the *Junos OS VPNs Library for Routing Devices*.

It is also possible to use the **ping** command to check the integrity of CCC LSPs. See [“Pinging CCC LSPs” on page 524](#) for more information.

Understanding Carrier-of-Carriers VPNs

IN THIS SECTION


- [Internet Service Provider as the Customer | 1646](#)
- [VPN Service Provider as the Customer | 1646](#)

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 1646](#)—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 1646](#)—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 126 on page 1645](#) illustrates the network architecture used for a carrier-of-carriers VPN service.

Figure 126: Carrier-of-Carriers VPN Architecture



8017197

This topic covers the following:

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 47 on page 1646](#).

Table 47: 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

Support for VPN service as the customer is supported on QFX10000 switches starting with Junos OS Release 17.1R1.

Release History Table

Release	Description
17.1R1	Support for VPN service as the customer is supported on QFX10000 switches starting with Junos OS Release 17.1R1.

RELATED DOCUMENTATION

[MPLS Feature Support on QFX Series and EX4600 Switches](#) | 26

Interprovider VPNs

Understanding 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.
- [“Understanding Carrier-of-Carriers VPNs” on page 1644](#)—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.

Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit (CLI Procedure)

You can configure an 802.1Q VLAN as an MPLS-based Layer 2 circuit on the switch to interconnect multiple customer sites with Layer 2 technology.

This topic describes configuring provider edge (PE) switches in an MPLS network using a circuit cross-connect (CCC) on a tagged VLAN interface (802.1Q VLAN) rather than a simple interface.

NOTE: You do not need to make any changes to existing provider switches in your MPLS network to support this type of configuration. For information on configuring provider switches, see [“Configuring MPLS on Provider Switches” on page 88](#).

NOTE: You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

NOTE: If you configure a physical interface as VLAN-tagged and with the `vlan-ccc` encapsulation, you cannot configure the associated logical interfaces with the `inet` family. Doing so could cause the logical interfaces to drop packets.

To configure a PE switch with a VLAN CCC and an MPLS-based Layer 2 circuit:

1. Configure OSPF (or IS-IS) on the loopback (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
```

2. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
```



```

user@switch# set interfaces lo0 unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address

```

4. Enable the MPLS protocol with CSPF disabled:

NOTE: CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. You need to disable CSPF for link protection to function properly on interarea paths.

```

[edit protocols]
user@switch# set mpls no-cspf

```

5. Configure the customer edge interface as a Layer 2 circuit from the local PE switch to the other PE switch:

```

[edit protocols]
user@switch# set l2circuit neighbor address interface interface-name virtual-circuit-id identifier

```

TIP: Use the switch address of the other switch as the neighbor address.

6. Configure MPLS on the core interfaces:

```

[edit protocols]
user@switch# set mpls interface interface-name
user@switch# set mpls interface interface-name
user@switch# set mpls interface interface-name

```

7. Configure LDP on the loopback interface and the core interfaces:

```

[edit protocols]
user@switch# set ldp interface lo0.0
user@switch# set ldp interface interface-name
user@switch# set ldp interface interface-name
user@switch# set ldp interface interface-name

```

8. Configure **family mpls** on the logical units of the core interfaces:


```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family mpls
user@switch# set interfaces interface-name unit logical-unit-number family mpls
user@switch# set interfaces interface-name unit logical-unit-number family mpls
```

NOTE: You can enable **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

9. Enable VLAN tagging on the customer edge interface of the local PE switch:

```
[edit]
user@switch# set interfaces interface-name vlan-tagging
```

10. Configure the customer edge interface to use VLAN CCC encapsulation:

```
[edit]
user@switch# set interfaces interface-name encapsulation vlan-ccc
```

11. Configure the logical unit of the customer edge interface with a VLAN ID:

NOTE: The VLAN ID cannot be configured on logical interface unit **0**. The logical unit number must be **1** or higher.

The same VLAN ID must be used when configuring the customer edge interface on the other PE switch.

```
[edit ]
user@switch# set interfaces interface-name logical-unit-number vlan-id vlan-id
```

When you have completed configuring one PE switch, follow the same procedures to configure the other PE switch.

NOTE: For EX Series switches, you must use the same type of switch for the other PE switch.
h.

RELATED DOCUMENTATION

[Configuring MPLS on Provider Switches | 88](#)

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

VLAN CCC Encapsulation on Transport Side of Pseudowire Client Logical Interfaces Overview

Currently, Junos OS does not allow the same VLAN ID to be configured on more than one logical interface under the same pseudowire client physical interface. To support **vlan-ccc** encapsulation on transport pseudowire service (PS) interface on the provider edge (PE) device, this restriction is removed and you can configure the same VLAN ID on more than one logical interface.

The primary reason to configure **vlan-ccc** on the transport PS interface is interoperability with the existing access and aggregate devices in the network. Currently, Junos OS supports **ethernet-ccc** encapsulation on the transport PS interface. Typically, while establishing a pseudowire connection, the access device initiates a VLAN-based pseudowire (also known as VLAN-tagged mode), and a PE router signals the Ethernet mode VLAN back to the access device. For this type of pseudowire connection to be established, you can use the **ignore-encapsulation-mismatch** statement. However, the Junos OS device (access device) might not support the **ignore-encapsulation-mismatch** statement and, as a result, the pseudowire connection is not formed. When the **ignore-encapsulation-mismatch** statement is not supported on the access device, you can configure **vlan-ccc** between the nodes to form a pseudowire connection.

The forwarding data path is not changed with the new **vlan-ccc** encapsulation on the transport PS interface and the behavior similar to that when the **ethernet-ccc** encapsulation is configured on the transport PS interface. The transport PS interface either encapsulates or de-encapsulate the outer Layer 2 header and MPLS headers on the transmitted or received packets on the WAN port. Inner Ethernet or VLAN headers of the packet are handled on pseudowire client service logical interfaces. You must configure pseudowire client service logical interfaces with appropriate VLAN IDs or VLAN tags.


The following sections provides details, along with a sample configuration, about pseudowire configuration from both access and aggregation nodes.

Pseudowire Configuration from Access Node

These pseudowires are set up using VLANs from the access node for customer devices attached to the Layer 2 circuit configured on access and PE routers with customer VLANs (C-VLANs). The ingress traffic (from the access node side) on the PE router is single VLAN tagged (inner Ethernet header), and thus the service logical interfaces must be configured with the same VLAN IDs corresponding to the C-VLAN IDs attached to the access node.

[Figure 127 on page 1652](#) provides the details of a transport PS interface from an access node (access node).

Figure 127: Pseudowire Client Transport Logical Interface from Access Node



The following example shows the configuration of a pseudowire client logical interface configuration on a PE router from an access node:

```

interfaces {
  ps0 {
    anchor-point lt-3;
    unit 0 {
      encapsulation VLAN-ccc;
      VLAN ID 100;
    }
    unit 1 {
      VLAN ID 100;
      family inet;
    }
  }
}

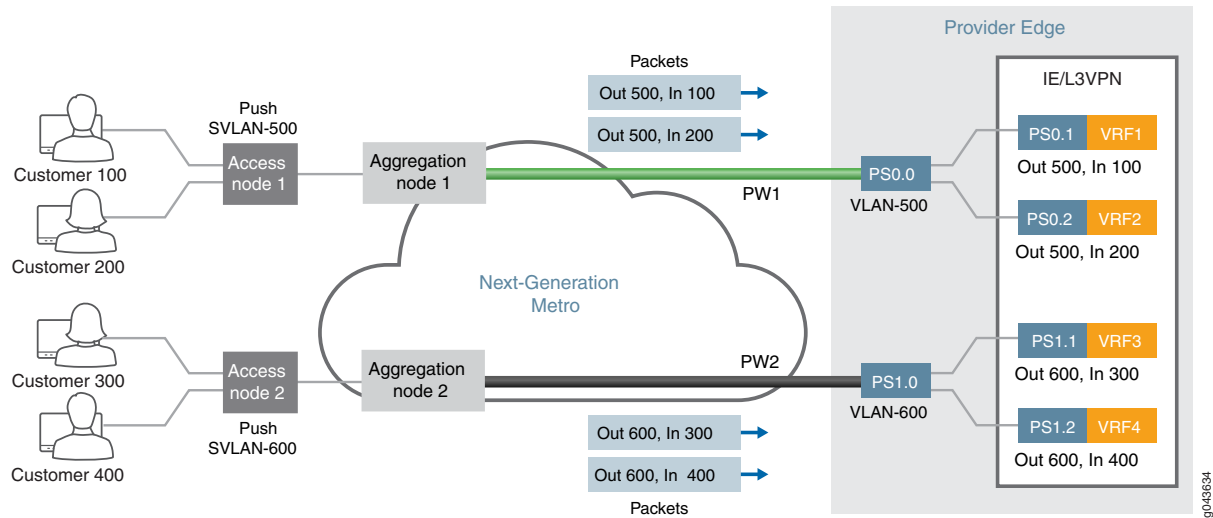
```

Pseudowire Configuration from Aggregation Node

In this case, the aggregation node processes a stacked VLAN (also known as Q-in-Q). The pseudowire originates from aggregation node and terminates on a PE router. The aggregation node pushes the service VLAN (S-VLAN) tag, and the PE router is expected to operate on two VLAN tags—the outer VLAN tag corresponds to an S-VLAN and the inner VLAN tag corresponds to a C-VLAN. The VLAN ID configured on the transport PS interface at the PE router must match the VLAN tag of the S-VLAN. On the pseudowire client service logical interface, the outer VLAN tag must be configured to match the S-VLAN and the inner VLAN tag must be configured to match the C-VLAN.

[Figure 128 on page 1653](#) provides the details of a transport PS interface from an aggregation node.

Figure 128: Pseudowire Client Transport Logical Interface from Aggregation Node



The following example shows the configuration of a pseudowire client logical interface configuration on a PE router from an aggregation node:

```

interfaces {
  ps0 {
    anchor-point lt-3;
    unit 0 {
      encapsulation VLAN-ccc;
      VLAN ID 500;
    }
    unit 1 {
      VLAN tags {
        outer 500;
        inner 100;
      }
    }
    unit 2 {
      VLAN tags {
        outer 500;
        inner 200;
      }
    }
  }
}

```


RELATED DOCUMENTATION

Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview

Pseudowire Subscriber Logical Interfaces Overview

Transmitting Nonstandard BPDUs

CCC protocol (and Layer 2 Circuit and Layer 2 VPN) configurations can transmit nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment. This is the default behavior on all supported PICs and requires no additional configuration.


The following PICs are supported on M320 and T Series routers:

- 1-port Gigabit Ethernet PIC
- 2-port Gigabit Ethernet PIC
- 4-port Gigabit Ethernet PIC
- 10-port Gigabit Ethernet PIC

TCC Overview

Translational cross-connect (TCC) is a switching concept that enables you to establish interconnections between a variety of Layer 2 protocols or circuits. It is similar to CCC. However, whereas CCC requires the same Layer 2 encapsulations on each side of a Juniper Networks router (such as PPP-to-PPP or Frame Relay-to-Frame Relay), TCC enables you to connect different types of Layer 2 protocols interchangeably. When you use TCC, combinations such as PPP-to-ATM (see [Figure 129 on page 1654](#)) and Ethernet-to-Frame Relay connections are possible.

Figure 129: TCC Example



The Layer 2 circuits and encapsulation types that can be interconnected by TCC are:

- Ethernet
- Extended VLANs
- PPP

- HDLC
- ATM
- Frame Relay

TCC works by removing the Layer 2 header when frames enter the router and adding a different Layer 2 header on the frames before they leave the router. In [Figure 129 on page 1654](#), the PPP encapsulation is stripped from the frames arriving at Router B, and the ATM encapsulation is added before the frames are sent to Router C.

Note that all control traffic is terminated at the interconnecting router (Router B). Examples of traffic controllers include the Link Control Protocol (LCP) and the Network Control Protocol (NCP) for PPP, keepalives for HDLC, and Local Management Interface (LMI) for Frame Relay.

TCC functionality is different from standard Layer 2 switching. TCC only swaps Layer 2 headers. No other processing, such as header checksums, TTL decrementing, or protocol handling is performed. TCC is supported for IPv4 only.

Address Resolution Protocol (ARP) packet policing on TCC Ethernet interfaces is effective for releases 10.4 and onwards.

You can configure TCC for interface switching and for Layer 2 VPNs. For more information about using TCC for virtual private networks (VPNs), see the *Junos OS VPNs Library for Routing Devices*.

Configuring Layer 2 Switching Cross-Connects Using CCC

IN THIS SECTION


- [Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects | 1656](#)
- [Configuring the CCC Connection for Layer 2 Switching Cross-Connects | 1661](#)
- [Configuring MPLS for Layer 2 Switching Cross-Connects | 1662](#)
- [Example: Configuring a Layer 2 Switching Cross-Connect | 1662](#)
- [Configuring Layer 2 Switching Cross-Connect on ACX5440 | 1664](#)

Layer 2 switching cross-connects join logical interfaces to form what is essentially Layer 2 switching. The interfaces that you connect must be of the same type.

Figure 130 on page 1656 illustrates a Layer 2 switching cross-connect. In this topology, Router A and Router C have Frame Relay connections to Router B, which is a Juniper Networks router. Circuit cross-connect (CCC) allows you to configure Router B to act as a Frame Relay (Layer 2) switch.

To configure Router B to act as a Frame Relay switch, you configure a circuit from Router A to Router C that passes through Router B, effectively configuring Router B as a Frame Relay switch with respect to these routers. This configuration allows Router B to transparently switch packets (frames) between Router A and Router C without regard to the packets' contents or the Layer 3 protocols. The only processing that Router B performs is to translate DLCI 600 to 750.

Figure 130: Layer 2 Switching Cross-Connect



If the Router A-to-Router B and Router B-to-Router C circuits were PPP, for example, the Link Control Protocol and Network Control Protocol exchanges occur between Router A and Router C. These messages are handled transparently by Router B, allowing Router A and Router C to use various PPP options (such as header or address compression and authentication) that Router B might not support. Similarly, Router A and Router C exchange keepalives, providing circuit-to-circuit connectivity status.

You can configure Layer 2 switching cross-connects on PPP, Cisco HDLC, Frame Relay, Ethernet, and ATM circuits. In a single cross-connect, only like interfaces can be connected.

To configure Layer 2 switching cross-connects, you must configure the following on the router that is acting as the switch (Router B in Figure 130 on page 1656):

Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects

IN THIS SECTION

- [Configuring ATM Encapsulation for Layer 2 Switching Cross-Connects | 1657](#)
- [Configuring Ethernet Encapsulation for Layer 2 Switching Cross-Connects | 1657](#)
- [Configuring Ethernet VLAN Encapsulation for Layer 2 Switching Cross-Connects | 1658](#)
- [Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching Cross-Connects | 1659](#)
- [Configuring Frame Relay Encapsulation for Layer 2 Switching Cross-Connects | 1660](#)
- [Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching Cross-Connects | 1661](#)

To configure Layer 2 switching cross-connects, configure the CCC encapsulation on the router that is acting as the switch (Router B in [Figure 130 on page 1656](#)).

NOTE: You cannot configure families on CCC interfaces; that is, you cannot include the **family** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

For instructions for configuring the encapsulation for Layer 2 switching cross-connects, see the following sections:

Configuring ATM Encapsulation for Layer 2 Switching Cross-Connects

For ATM circuits, specify the encapsulation when configuring the virtual circuit (VC). Configure each VC as a circuit or a regular logical interface by including the following statements:

```
at-fpc/pic/port {
  atm-options {
    vpi vpi-identifier maximum-vcs maximum-vcs;
  }
  unit logical-unit-number {
    encapsulation encapsulation-type;
    point-to-point; # Default interface type
    vci vpi-identifier.vci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

Configuring Ethernet Encapsulation for Layer 2 Switching Cross-Connects

For Ethernet circuits, specify **ethernet-ccc** in the **encapsulation** statement. This statement configures the entire physical device. For these circuits to work, you must also configure a logical interface (unit 0).

Ethernet interfaces with standard Tag Protocol Identifier (TPID) tagging can use Ethernet CCC encapsulation. On M Series Multiservice Edge Routers, except the M320, one-port Gigabit Ethernet, two-port Gigabit Ethernet, four-port Gigabit Ethernet, and four-port Fast Ethernet PICs can use Ethernet CCC encapsulation. On T Series Core Routers and M320 routers, one-port Gigabit Ethernet and two-port Gigabit Ethernet PICs installed in FPC2 can use Ethernet CCC encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.

```
fe-fpc/pic/port {
```



```
encapsulation ethernet-ccc;
unit 0;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Configuring Ethernet VLAN Encapsulation for Layer 2 Switching Cross-Connects

An Ethernet virtual LAN (VLAN) circuit can be configured using either the **vlan-ccc** or **extended-vlan-ccc** encapsulation. If you configure the **extended-vlan-ccc** encapsulation on the physical interface, you cannot configure the **inet** family on the logical interfaces. Only the **ccc** family is allowed. If you configure the **vlan-ccc** encapsulation on the physical interface, both the **inet** and **ccc** families are supported on the logical interfaces. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For encapsulation type **vlan-ccc**, VLAN IDs from 512 through 4094 are reserved for CCC VLANs. For the **extended-vlan-ccc** encapsulation type, all VLAN IDs 1 and higher are valid. VLAN ID 0 is reserved for tagging the priority of frames.

NOTE: Some vendors use the proprietary TPIDs 0x9100 and 0x9901 to encapsulate a VLAN-tagged packet into a VLAN-CCC tunnel to interconnect a geographically separated metro Ethernet network. By configuring the **extended-vlan-ccc** encapsulation type, a Juniper Networks router can accept all three TPIDs (0x8100, 0x9100, and 0x9901).

Configure an Ethernet VLAN circuit with the **vlan-ccc** encapsulation as follows:

```
interfaces {
  type-fpc/pic/port {
    vlan-tagging;
    encapsulation vlan-ccc;
    unit logical-unit-number {
      encapsulation vlan-ccc;
      vlan-id vlan-id;
    }
  }
}
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces]

- [edit logical-systems *logical-system-name* interfaces]

Configure an Ethernet VLAN circuit with the **extended-vlan-ccc** encapsulation statement as follows:

```
interfaces {
  type-fpc/pic/port {
    vlan-tagging;
    encapsulation extended-vlan-ccc;
    unit logical-unit-number {
      vlan-id vlan-id;
      family ccc;
    }
  }
}
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Whether you configure the encapsulation as **vlan-ccc** or **extended-vlan-ccc**, you must enable VLAN tagging by including the **vlan-tagging** statement.

Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching Cross-Connects

You can configure aggregated Ethernet interfaces for CCC connections and for Layer 2 virtual private networks (VPNs).

Aggregated Ethernet interfaces configured with VLAN tagging can be configured with multiple logical interfaces. The only encapsulation available for aggregated Ethernet logical interfaces is **vlan-ccc**. When you configure the **vlan-id** statement, you are limited to VLAN IDs 512 through 4094.

Aggregated Ethernet interfaces configured without VLAN tagging can be configured only with the **ethernet-ccc** encapsulation. All untagged Ethernet packets received are forwarded based on the CCC parameters.

To configure aggregated Ethernet interfaces for CCC connections, include the **ae0** statement at the [edit interfaces] hierarchy level:

```
[edit interfaces]
ae0 {
  encapsulation (ethernet-ccc | extended-vlan-ccc | vlan-ccc);
  vlan-tagging;
  aggregated-ether-options {
    minimum-links links;
    link-speed speed;
```



```

    }
    unit logical-unit-number {
        encapsulation vlan-ccc;
        vlan-id identifier;
        family ccc;
    }
}

```

Be aware of the following limitations when configuring CCC connections over aggregated Ethernet interfaces:

- If you configured load balancing between child links, be aware that a different hash key is used to distribute packets among the child links. Standard aggregated interfaces have family inet configured. An IP version 4 (IPv4) hash key (based on the Layer 3 information) is used to distribute packets among the child links. A CCC connection over an aggregated Ethernet interface has family ccc configured instead. Instead of an IPv4 hash key, an MPLS hash key (based on the destination media access control [MAC] address) is used to distributed packets among the child links.
- The extended-vlan-ccc encapsulation is not supported on the 12-port Fast Ethernet PIC and the 48-port Fast Ethernet PIC.
- The Junos OS does not support the Link Aggregation Control Protocol (LACP) when an aggregated interface is configured as a VLAN (with vlan-ccc encapsulation). LACP can be configured only when the aggregated interface is configured with the ethernet-ccc encapsulation.

For more information about how to configure aggregated Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

Configuring Frame Relay Encapsulation for Layer 2 Switching Cross-Connects

For Frame Relay circuits, specify the encapsulation when configuring the DLCI. Configure each DLCI as a circuit or a regular logical interface. The DLCI for regular interfaces must be from 1 through 511. For CCC interfaces, it must be from 512 through 4094.

```

interfaces {
    type-fpc/pic/port {
        unit logical-unit-number {
            dlcid dlcid-identifier;
            encapsulation encapsulation-type;
            point-to-point; # Default interface type
        }
    }
}

```

You can configure these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching Cross-Connects

For PPP and Cisco HDLC circuits, specify the encapsulation in the **encapsulation** statement. This statement configures the entire physical device. For these circuits to work, you must configure a logical interface (unit 0).

```
interfaces type-fpc/pic/port {
  encapsulation encapsulation-type;
  unit 0;
}
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces *type-fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces *type-fpc/pic/port*]

Configuring the CCC Connection for Layer 2 Switching Cross-Connects

To configure Layer 2 switching cross-connects, define the connection between the two circuits by including the **interface-switch** statement. You configure this connection on the router that is acting as the switch (Router B in [Figure 130 on page 1656](#)). The connection joins the interface that comes from the circuit's source to the interface that leads to the circuit's destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.

```
interface-switch connection-name {
  interface interface-name.unit-number;
  interface interface-name.unit-number;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

Configuring MPLS for Layer 2 Switching Cross-Connects

For Layer 2 switching cross-connects to work, you must enable MPLS on the router by including at least the following statements. This minimum configuration enables MPLS on a logical interface for the switching cross-connect.

Include the **family mpls** statement:

```
family mpls;
```

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You can then specify this logical interface in the MPLS protocol configuration:

```
mpls {
  interface interface-name; # Required to enable MPLS on the interface
}
```


You can configure these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

Example: Configuring a Layer 2 Switching Cross-Connect

Configure a full-duplex Layer 2 switching cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the virtual switch. See the topology in [Figure 131 on page 1662](#) and [Figure 132 on page 1663](#).

Figure 131: Topology of a Frame Relay Layer 2 Switching Cross-Connect



```
[edit]
interfaces {
  so-1/0/0 {
```




```

encapsulation frame-relay-ccc;
unit 1 {
    point-to-point;
    encapsulation frame-relay-ccc;
    dlci 600;
}
}
so-2/0/0 {
    encapsulation frame-relay-ccc;
    unit 2 {
        point-to-point;
        encapsulation frame-relay-ccc;
        dlci 750;
    }
}
}
protocols {
    connections {
        interface-switch router-a-to-router-c {
            interface so-1/0/0.1;
            interface so-2/0/0.2;
        }
    }
    mpls {
        interface all;
    }
}
}

```

Figure 132: Sample Topology of a VLAN Layer 2 Switching Cross-Connect



```

[edit]
interfaces {
    ge-2/1/0 {
        vlan-tagging;
        encapsulation vlan-ccc;
        unit 0 {
            encapsulation vlan-ccc;
            vlan-id 600;
        }
    }
}

```



```

    }
  }
  ge-2/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    unit 0 {
      encapsulation vlan-ccc;
      vlan-id 600;
    }
    unit 1 {
      family inet {
        vlan-id 1;
        address 10.9.200.1/24;
      }
    }
  }
}
protocols {
  mpls {
    interface all;
  }
  connections {
    interface-switch layer2-sw {
      interface ge-2/1/0.0;
      interface ge-2/2/0.0;
    }
  }
}

```

Configuring Layer 2 Switching Cross-Connect on ACX5440

Starting in Junos OS Release 19.3R1, you can leverage the hardware support available for cross-connects on the ACX5448 device with the Layer 2 local switching functionality using certain models. With this support, you can provide the EVP and Ethernet Virtual Private Line (EVPL) services..

Local-switching with the following forwarding models are supported:

- VLAN-CCC (logical interface-level local-switching) without any map.
- VLAN-CCC (logical interface-level local-switching) with the following vlan-maps:
 - Push 0x8100.pushVLAN (QinQ type)
 - Swap 0x8100.swapVLAN
- Aggregated Ethernet (AE) static interfaces.

- AE interfaces with LACP, load-balance all active mode.
- Local-switching end-interface support for AE or LAG interface (one non-AE interface and other AE interface).
- Local-switching both interface as AE or LAG interfaces.

To enable Layer 2 local switching on the ACX5448 device, you can use the existing configuration statements for Layer 2 circuits. For example,

```
[edit protocols l2circuit]
local-switching {
  interface interface1 {
    end-interface interface3;
    ignore-encapsulation-mismatch;
    ignore-mtu-mismatch;
  }
}
```

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, you can leverage the hardware support available for cross-connects on the ACX5448 device with the Layer 2 local switching functionality using certain models. With this support, you can provide the EVP and Ethernet Virtual Private Line (EVPL) services.

RELATED DOCUMENTATION

Configuring CCC Encapsulation for Layer 2 VPNs
Understanding Encapsulation on an Interface


Configuring MPLS LSP Tunnel Cross-Connects Using CCC

IN THIS SECTION

- [Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects | 1667](#)
- [Configuring the CCC Connection for LSP Tunnel Cross-Connects | 1669](#)
- [Example: Configuring an LSP Tunnel Cross-Connect | 1669](#)

MPLS tunnel cross-connects between interfaces and LSPs allow you to connect two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the conduit. The topology in [Figure 133 on page 1666](#) illustrates an MPLS LSP tunnel cross-connect. In this topology, two separate networks, in this case ATM access networks, are connected through an IP backbone. CCC allows you to establish an LSP tunnel between the two domains. With LSP tunneling, you tunnel the ATM traffic from one network across a SONET backbone to the second network by using an MPLS LSP.

Figure 133: MPLS Tunnel Cross-Connect



When traffic from Router A (VC 234) reaches Router B, it is encapsulated and placed into an LSP, which is sent through the backbone to Router C. At Router C, the label is removed, and the packets are placed onto the ATM permanent virtual circuit (PVC) (VC 591) and sent to Router D. Similarly, traffic from Router D (VC 591) is sent over an LSP to Router B, then placed on VC 234 to Router A.

You can configure LSP tunnel cross-connect on PPP, Cisco HDLC, Frame Relay, and ATM circuits. In a single cross-connect, only like interfaces can be connected.

When you use MPLS tunnel cross-connects to support IS-IS, you must ensure that the LSP's maximum transmission unit (MTU) can, at a minimum, accommodate a 1492-octet IS-IS protocol data unit (PDU) in addition to the link-level overhead associated with the technology being connected.

For the tunnel cross-connects to work, the IS-IS frame size on the edge routers (Routers A and D in [Figure 134 on page 1669](#)) must be smaller than the LSP's MTU.

NOTE: Frame size values do not include the frame check sequence (FCS) or delimiting flags.

To determine the LSP MTU required to support IS-IS, use the following calculation:

IS-IS MTU (minimum 1492, default 1497) + frame overhead + 4 (MPLS shim header) = Minimum LSP MTU

The framing overhead varies based on the encapsulation being used. The following lists the IS-IS encapsulation overhead values for various encapsulations:

- ATM
 - AAL5 multiplex—8 bytes (RFC 1483)

- VC multiplex—0 bytes
- Frame Relay
 - Multiprotocol—2 bytes (RFCs 1490 and 2427)
 - VC multiplex—0 bytes
- HDLC—4 bytes
- PPP—4 bytes
- VLAN—21 bytes (802.3/LLC)

For IS-IS to work over VLAN-CCC, the LSP's MTU must be at least 1513 bytes (or 1518 for 1497-byte PDUs). If you increase the size of a Fast Ethernet MTU above the default of 1500 bytes, you might need to explicitly configure jumbo frames on intervening equipment.

To modify the MTU, include the **mtu** statement when configuring the logical interface family at the **[edit interfaces interface-name unit logical-unit-number encapsulation family]** hierarchy level. For more information about setting the MTU, see the *Junos OS Network Interfaces Library for Routing Devices*.

To configure an LSP tunnel cross-connect, you must configure the following on the interdomain router (Router B in [Figure 134 on page 1669](#)):

Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects

To configure LSP tunnel cross-connects, you must configure the CCC encapsulation on the ingress and egress routers (Router B and Router C, respectively, in [Figure 134 on page 1669](#)).

NOTE: You cannot configure families on CCC interfaces; that is, you cannot include the **family** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level.

For PPP or Cisco HDLC circuits, include the **encapsulation** statement to configure the entire physical device. For these circuits to work, you must configure logical unit 0 on the interface.

```
type-fpc/pic/port {
  encapsulation (ppp-ccc | cisco-hdlc-ccc);
  unit 0;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems logical-system-name interfaces]**

For ATM circuits, specify the encapsulation when configuring the VC by including the following statements. For each VC, you configure whether it is a circuit or a regular logical interface.

```
at-fpc/pic/port {
  atm-options {
    vpi vpi-identifier maximum-vcs maximum-vcs;
  }
  unit logical-unit-number {
    point-to-point; # Default interface type
    encapsulation atm-ccc-vc-mux;
    vci vpi-identifier.vci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

For Frame Relay circuits, include the following statements to specify the encapsulation when configuring the DLCI. For each DLCI, you configure whether it is a circuit or a regular logical interface. The DLCI for regular interfaces must be in the range 1 through 511. For CCC interfaces, it must be in the range 512 through 1022.

```
type-fpc/pic/port {
  encapsulation frame-relay-ccc;
  unit logical-unit-number {
    point-to-point; # default interface type
    encapsulation frame-relay-ccc;
    dlci dlci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

For more information about the **encapsulation** statement, see the *Junos OS Network Interfaces Library for Routing Devices*.

Configuring the CCC Connection for LSP Tunnel Cross-Connects

To configure LSP tunnel cross-connects, include the **remote-interface-switch** statement to define the connection between the two circuits on the ingress and egress routers (Router B and Router C, respectively, in [Figure 134 on page 1669](#)). The connection joins the interface or LSP that comes from the circuit's source to the interface or LSP that leads to the circuit's destination. When you specify the interface name, include the logical portion of the name, which corresponds to the logical unit number. For the cross-connect to be bidirectional, you must configure cross-connects on two routers.

```
remote-interface-switch connection-name {
  interface interface-name.unit-number;
  transmit-lsp label-switched-path;
  receive-lsp label-switched-path;
}
```


You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

Example: Configuring an LSP Tunnel Cross-Connect

Configure a full-duplex MPLS LSP tunnel cross-connect from Router A to Router D, passing through Router B and Router C. See the topology in [Figure 134 on page 1669](#).

Figure 134: Example Topology of MPLS LSP Tunnel Cross-Connect



On Router B:

```
[edit]
interfaces {
  at-7/1/1 {
    atm-options {
      vpi 1 maximum-vcs 600;
    }
    unit 1 {
```



```

        point-to-point; # default interface type
        encapsulation atm-ccc-vc-mux;
        vci 1.234;
    }
}
}
protocols {
    connections {
        remote-interface-switch router-b-to-router-c {
            interface at-7/1/1.1;
            transmit-lsp lsp1;
            receive-lsp lsp2;
        }
    }
}

```

On Router C:

```

[edit]
interfaces {
    at-3/0/0 {
        atm-options {
            vpi 2 maximum-vcs 600;
        }
        unit 2 {
            point-to-point; # default interface type
            encapsulation atm-ccc-vc-mux;
            vci 2.591;
        }
    }
}
protocols {
    connections {
        remote-interface-switch router-b-to-router-c {
            interface at-3/0/0.2;
            transmit-lsp lsp2;
            receive-lsp lsp1;
        }
    }
}

```


Configuring TCC

IN THIS SECTION

- [Configuring the Encapsulation for Layer 2 Switching TCCs | 1671](#)
- [Configuring the Connection for Layer 2 Switching TCCs | 1675](#)
- [Configuring MPLS for Layer 2 Switching TCCs | 1676](#)

This section describes how to configure translational cross-connect (TCC).

To configure TCC, you must perform the following tasks on the router that is acting as the switch:

Configuring the Encapsulation for Layer 2 Switching TCCs

IN THIS SECTION

- [Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching TCCs | 1672](#)
- [Configuring ATM Encapsulation for Layer 2 Switching TCCs | 1672](#)
- [Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs | 1672](#)
- [Configuring Ethernet Encapsulation for Layer 2 Switching TCCs | 1673](#)
- [Configuring Ethernet Extended VLAN Encapsulation for Layer 2 Switching TCCs | 1674](#)
- [Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations | 1675](#)

To configure a Layer 2 switching TCC, specify the TCC encapsulation on the desired interfaces of the router that is acting as the switch.

NOTE: You cannot configure standard protocol families on TCC or CCC interfaces. Only the CCC family is allowed on CCC interfaces, and only the TCC family is allowed on TCC interfaces.

For Ethernet circuits and Ethernet extended VLAN circuits, you must also configure the Address Resolution Protocol (ARP). See [“Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations” on page 1675](#).

Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching TCCs

For PPP and Cisco HDLC circuits, configure the encapsulation type for the entire physical device by specifying the appropriate value for the **encapsulation** statement. For these circuits to work, you must also configure the logical interface **unit 0**.

```
encapsulation (ppp-tcc | cisco-hdlc-tcc);
unit 0{...}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

Configuring ATM Encapsulation for Layer 2 Switching TCCs

For ATM circuits, configure the encapsulation type by specifying the appropriate value for the **encapsulation** statement in the virtual circuit (VC) configuration. Specify whether each VC is a circuit or a regular logical interface.

```
atm-options {
  vpi vpi-identifier maximum-vcs maximum-vcs;
}
unit logical-unit-number {
  encapsulation (atm-tcc-vc-mux | atm-tcc-snap);
  point-to-point;
  vci vpi-identifier.vci-identifier;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *at-fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces *at-fpc/pic/port*]

Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs

For Frame Relay circuits, configure the encapsulation type by specifying the value **frame-relay-tcc** for the **encapsulation** statement when configuring the data-link connection identifier (DLCI). You configure each DLCI as a circuit or a regular logical interface. The DLCI for regular interfaces must be in the range from 1 through 511, but for TCC and CCC interfaces it must be in the range from 512 through 1022.

```
encapsulation frame-relay-tcc;
unit logical-unit-number {
  dlci dlci-identifier;
  encapsulation frame-relay-tcc;
```



```
point-to-point;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

Configuring Ethernet Encapsulation for Layer 2 Switching TCCs

For Ethernet TCC circuits, configuring the encapsulation type for the entire physical device by specifying the value **ethernet-tcc** for the **encapsulation** statement.

You must also specify static values for a remote address and a proxy address at the [edit interfaces *interface-name* unit *unit-number* family *tcc*] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family *tcc*] hierarchy level.

The remote address is associated with the TCC switching router's Ethernet neighbor; in the **remote** statement you must specify both the IP address and the media access control (MAC) address of the Ethernet neighbor. The proxy address is associated with the TCC router's other neighbor connected by the unlike link; in the **proxy** statement you must specify the IP address of the non-Ethernet neighbor.

You can configure Ethernet TCC encapsulation for the interfaces on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Fast Ethernet, and 4-port Gigabit Ethernet PICs.

```
encapsulation ethernet-tcc;
unit logical-unit-number {
  family tcc {
    proxy {
      inet-address ip-address;
    }
    remote {
      inet-address ip-address;
      mac-address mac-address;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces (fe | ge)-*fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces (fe | ge)-*fpc/pic/port*]

NOTE: For Ethernet circuits, you must also configure the Address Resolution Protocol (ARP). See [“Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations” on page 1675](#).

Configuring Ethernet Extended VLAN Encapsulation for Layer 2 Switching TCCs

For Ethernet extended VLAN circuits, configure the encapsulation type for the entire physical device by specifying the value **extended-vlan-tcc** for the **encapsulation** statement.

You must also enable VLAN tagging. Ethernet interfaces in VLAN mode can have multiple logical interfaces. With encapsulation type **extended-vlan-tcc**, all VLAN IDs from 0 through 4094 are valid, up to a maximum of 1024 VLANs. As with Ethernet circuits, you must also specify a proxy address and a remote address at the **[edit interfaces *interface-name* unit *logical-unit-number* family tcc]** or **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family tcc]** hierarchy level (see [“Configuring Ethernet Encapsulation for Layer 2 Switching TCCs” on page 1673](#)).

```
encapsulation extended-vlan-tcc;
vlan-tagging;
unit logical-unit-number {
    vlan-id identifier;
    family tcc;
    proxy {
        inet-address ip-address;
    }
    remote {
        inet-address ip-address;
        mac-address mac-address;
    }
}
```

You can configure these statements at the following hierarchy levels:

- **[edit interfaces *interface-name*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name*]**

NOTE: For Ethernet extended VLAN circuits, you must also configure the Address Resolution Protocol (ARP). See [“Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations” on page 1675](#).

Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations

For Ethernet and Ethernet extended VLAN circuits with TCC encapsulation, you must also configure ARP. Because TCC simply removes one Layer 2 header and adds another, the default form of dynamic ARP is not supported; you must configure static ARP.

Because remote and proxy addresses are specified on the router performing TCC switching, you must apply the static ARP statement to the Ethernet-type interfaces of the routers that connect to the TCC-switched router. The **arp** statement must specify the IP address and the MAC address of the remotely connected neighbor by use of the unlike Layer 2 protocol on the far side of the TCC switching router.

```
arp ip-address mac mac-address;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *ip-address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *ip-address*]

Configuring the Connection for Layer 2 Switching TCCs

You must configure the connection between the two circuits of the Layer 2 switching TCC on the router acting as the switch. The connection joins the interface coming from the circuit's source to the interface leading to the circuit's destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted from the second interface, and those received on the second interface are transmitted from the first.

To configure a connection for a local interface switch, include the following statements:

```
interface-switch connection-name {
    interface interface-name.unit-number;
}
lsp-switch connection-name {
    transmit-lsp lsp-number;
    receive-lsp lsp-number;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

To configure a connection for a remote interface switch, include the following statements:


```
remote-interface-switch connection-name {
    interface interface-name.unit-number;
    interface interface-name.unit-number;
    transmit-lsp lsp-number;
    receive-lsp lsp-number;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

Configuring MPLS for Layer 2 Switching TCCs

For a Layer 2 switching TCC to work, you must enable MPLS on the router by including at least the following statements. This minimum configuration enables MPLS on a logical interface for the switching cross-connect.

Include the **family mpls** statement:

```
family mpls;
```

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You can then specify this logical interface in the MPLS protocol configuration:

```
mpls {
    interface interface-name; # Required to enable MPLS on the interface
}
```

You can configure these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

NOTE: MPLS LSP link protection does not support TCC.

RELATED DOCUMENTATION

[TCC Overview | 1654](#)
[Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits](#)


CCC and TCC Graceful Restart

CCC and TCC graceful restart allows Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the **remote-interface-switch** or **lsp-switch** statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the PE routers and P routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

[Figure 135 on page 1677](#) illustrates how graceful restart might work on a CCC connection between two CE routers.

Figure 135: Remote Interface Switch Connecting Two CE Routers Using CCC



PE Router A is the ingress for the transmit LSP from PE Router A to PE Router B and the egress for the receive LSP from PE Router B to PE Router A. With RSVP graceful restart enabled on all the PE and P routers, the following occurs when PE router A restarts:

- PE Router A preserves the forwarding state associated with the CCC routes (those from CCC to MPLS and from MPLS to CCC).
- Traffic flows without disruption from CE router to CE router.
- After the restart, PE Router A preserves the label for the LSP for which PE Router A is the egress (the receive LSP, for example). The transmit LSP from PE Router A to PE Router B can derive new label mappings, but should not cause any traffic disruption.

Configuring CCC and TCC Graceful Restart

To enable CCC and TCC graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library*.

CCC and TCC graceful restart depend on RSVP graceful restart. If you disable RSVP graceful restart, CCC and TCC graceful restart will not work. For more information about RSVP graceful restart, see [“RSVP Graceful Restart” on page 881](#) and [“Configuring RSVP Graceful Restart” on page 883](#).

Configuring an MPLS-Based VLAN CCC Using the Connection Method (CLI Procedure)

You can configure an 802.1Q VLAN as an MPLS-based connection using EX8200 and EX4500 switches to interconnect multiple customer sites with Layer 2 technology.

This topic describes configuring provider edge (PE) switches in an MPLS network using a circuit cross-connect (CCC) on a tagged VLAN interface (802.1Q VLAN) rather than a simple interface.

NOTE: You do not need to make any changes to existing provider switches in your MPLS network to support this type of configuration. For information on configuring provider switches, see [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

NOTE: You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

NOTE: If you configure a physical interface as VLAN-tagged and with the `vlan-ccc` encapsulation, you cannot configure the associated logical interfaces with the `inet` family. Doing so could cause the logical interfaces to drop packets.

To configure a PE switch with a VLAN CCC and an MPLS-based connections:

1. Configure OSPF (or IS-IS) on the loopback (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
```

2. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
```



```

user@switch# set interfaces lo0 unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address

```

4. Enable the MPLS protocol with **cspf** disabled:

NOTE: CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. You need to disable CSPF for link protection to function properly on interarea paths.

```

[edit protocols]
user@switch# set mpls no-cspf

```

5. Enable VLAN tagging on the customer edge interface of the local PE switch:

```

[edit]
user@switch# set interfaces interface-name vlan-tagging

```

6. Configure the customer edge interface to use encapsulation **vlan-ccc**:

```

[edit]
user@switch# set interfaces interface-name encapsulation vlan-ccc

```

7. Configure the logical unit of the customer edge interface with a VLAN ID:

NOTE: The VLAN ID cannot be configured on logical interface unit **0**.

The same VLAN ID must be used when configuring the customer edge interface on the other PE switch.

```

[edit ]
user@switch# set interfaces interface-name logical-unit-number vlan-id

```

8. Define the label switched path (LSP):

```

[edit protocols]
user@switch# set mpls label-switched-path lsp-name from address
user@switch# set mpls label-switched-path lsp-name to address

```


TIP: You will need to use the specified LSP name again when configuring the CCC.

9. Configure the connection between the two circuits in the CCC connection

```
[edit protocols]
user@switch# set connections remote-interface-switch interface-switch interface local-interface
user@switch# set connections remote-interface-switch interface-switch transmit-lsp destination-lsp
user@switch# set connections remote-interface-switch interface-switch receive-lsp source-lsp
```

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

Configuring CCC Switching for Point-to-Multipoint LSPs

IN THIS SECTION

- [Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers | 1682](#)
- [Configuring Local Receivers on a Point-to-Multipoint CCC LSP Switch on Ingress PE Routers | 1682](#)
- [Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers | 1683](#)

You can configure circuit cross-connect (CCC) between two circuits to switch traffic from interfaces to point-to-multipoint LSPs. This feature is useful for handling multicast or broadcast traffic (for example, a digital video stream).

To configure CCC switching for point-to-multipoint LSPs, you do the following:

- On the ingress provider edge (PE) router, you configure CCC to switch traffic from an incoming interface to a point-to-multipoint LSP.
- On the egress PE, you configure CCC to switch traffic from an incoming point-to-multipoint LSP to an outgoing interface.

The CCC connection for point-to-multipoint LSPs is unidirectional.

For more information about point-to-multipoint LSPs, see [“Point-to-Multipoint LSPs Overview” on page 691](#).

To configure a CCC connection for a point-to-multipoint LSP, complete the steps in the following sections:

Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers

To configure the ingress PE router with a CCC switch for a point-to-multipoint LSP, include the **p2mp-transmit-switch** statement:

```
p2mp-transmit-switch switch-name {
    input-interface input-interface-name.unit-number;
    transmit-p2mp-lsp transmitting-lsp;
}
```

You can include the **p2mp-transmit-switch** statement at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

switch-name specifies the name of the ingress CCC switch.

input-interface *input-interface-name.unit-number* specifies the name of the ingress interface.

transmit-p2mp-lsp *transmitting-lsp* specifies the name of the transmitting point-to-multipoint LSP.

Configuring Local Receivers on a Point-to-Multipoint CCC LSP Switch on Ingress PE Routers

In addition to configuring an incoming CCC interface to a point-to-multipoint LSP on an ingress PE router, you can also configure CCC to switch traffic on an incoming CCC interface to one or more outgoing CCC interfaces by configuring output interfaces as local receivers.

To configure output interfaces, include the **output-interface** statement at the [edit protocols connections **p2mp-transmit-switch** *p2mp-transmit-switch-name*] hierarchy level.

```
[edit protocols connections]
p2mp-transmit-switch pc-ccc {
    input-interface fe-1/3/1.0;
    transmit-p2mp-lsp myp2mp;
    output-interface [fe-1/3/2.0 fe-1/3/3.0];
}
```

You can configure one or more output interfaces as local receivers on the ingress PE router using this statement.

Use the **show connections p2mp-transmit-switch** (*extensive* | *history* | *status*), **show route ccc** *<interface-name>* (*detail* | *extensive*), and **show route forwarding-table ccc** *<interface-name>* (*detail* | *extensive*) commands to view details of the local receiving interfaces on the ingress PE router.

Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers

To configure the CCC switch for a point-to-multipoint LSP on the egress PE router, include the **p2mp-receive-switch** statement.

```
p2mp-receive-switch switch-name {
    output-interface [ output-interface-name.unit-number ];
    receive-p2mp-lsp receptive-lsp;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

switch-name specifies the name of the egress CCC switch.

output-interface [*output-interface-name.unit-number*] specifies the name of one or more egress interfaces.

receive-p2mp-lsp *receptive-lsp* specifies the name of the receptive point-to-multipoint LSP.

Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN (CLI Procedure)

You can configure an 802.1Q VLAN as an MPLS-based Layer 2 virtual private network (VPN) using EX8200 and EX4500 switches to interconnect multiple customer sites with Layer 2 technology.

This topic describes configuring provider edge (PE) switches in an MPLS network using a circuit cross-connect (CCC) on a tagged VLAN interface (802.1Q VLAN) rather than a simple interface.

NOTE: You do not need to make any changes to existing provider switches in your MPLS network to support this type of configuration. For information on configuring provider switches, see [“Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\)” on page 99](#).

NOTE: You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

NOTE: If you configure a physical interface as VLAN-tagged and with the `vlan-ccc` encapsulation, you cannot configure the associated logical interfaces with the `inet` family. Doing so could cause the logical interfaces to drop packets.

To configure a PE switch with a VLAN CCC and an MPLS-based Layer 2 VPN:

1. Configure OSPF (or IS-IS) on the loopback (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
```

2. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit logical-unit-number family inet address address
```



```

user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address

```

4. Enable the MPLS protocol with **cspf** disabled:

NOTE: CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. You need to disable CSPF for link protection to function properly on interarea paths.

```

[edit protocols]
user@switch# set mpls no-cspf

```

5. Define the label switched path (LSP):

```

[edit protocols]
user@switch# set mpls label-switched-path lsp_name to address

```

TIP: You will need to use the specified LSP name again when configuring the CCC.

6. Configure MPLS on the core interfaces:

```

[edit protocols]
user@switch# set mpls interface interface-name
user@switch# set mpls interface interface-name
user@switch# set mpls interface interface-name

```

7. Configure RSVP on the loopback interface and the core interfaces:

```

[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface interface-name
user@switch# set rsvp interface interface-name
user@switch# set rsvp interface interface-name

```

8. Configure **family mpls** on the logical units of the core interfaces:

```

[edit]

```



```

user@switch# set interfaces interface-name unit logical-unit-number family mpls
user@switch# set interfaces interface-name unit logical-unit-number family mpls
user@switch# set interfaces interface-name unit logical-unit-number family mpls

```

NOTE: You can enable **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

9. Enable VLAN tagging on the customer edge interface of the local PE switch:

```

[edit]
user@switch# set interfaces interface-name vlan-tagging

```

10. Configure the customer edge interface to use encapsulation **vlan-ccc**:

```

[edit]
user@switch# set interfaces interface-name encapsulation vlan-ccc

```

11. Configure the logical unit of the customer edge interface with a VLAN ID:

NOTE: The VLAN ID cannot be configured on logical interface unit **0**. The logical unit number must be **1** or higher.

The same VLAN ID must be used when configuring the customer edge interface on the other PE switch.

```

[edit ]
user@switch# set interfaces interface-name logical-unit-number vlan-id vlan-id

```

12. Configure BGP, specifying the loopback address as the local address and enabling **family l2vpn signaling**:

```

[edit protocols bgp]
user@switchPE1# set local-address address family l2vpn signaling

```

13. Configure the BGP group, specifying the group name and type:

```

[edit protocols bgp]
user@switchPE1# set group ibgp type internal

```

14. 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 address
```

15. Configure the routing instance, specifying the routing-instance name and using **l2vpn** as the instance type:

```
[edit routing-instances]
user@switchPE1# set routing-instance-name instance-type l2vpn
```

16. Configure the routing instance to apply to the customer edge interface:

```
[edit routing-instances]
user@switchPE1# set routing-instance-name interface interface-name
```

17. Configure the routing instance to use a route distinguisher:

```
[edit routing-instances]
user@switchPE1# set routing-instance-name route-distinguisher address
```

18. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@switchPE1# set routing-instance-name vrf-target community
```

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 Configuration Guide](#).

19. Configure the protocols and encapsulation type used by the routing instance:

```
[edit routing-instances]
user@switchPE1# set routing-instance-name protocols l2vpn encapsulation-type ethernet-vlan
```

20. Apply the routing instance to a customer edge interface and specify a description for it:

```
[edit routing-instances]
user@switchPE1# set routing-instance-name protocols interface interface-name description description
```

21. Configure the routing-instance protocols site:

```
[edit routing-instances]
```



```
user@switchPE1# set routing-instance-name protocols l2vpn site site-name site-identifier identifier  
remote-site-id identifier
```

NOTE: The remote site ID (configured with the **remote-site-id** statement) corresponds to the site ID (configured with the **site-identifier** statement) configured on the other PE switch.

When you have completed configuring one PE switch, follow the same procedures to configure the other PE switch.

NOTE: You must use the same type of switch for the other PE switch. You cannot use an EX8200 as one PE switch and use an EX3200 or EX4200 as the other PE switch.

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit (CLI Procedure)

You can configure an 802.1Q VLAN as an MPLS-based Layer 2 circuit on the switch to interconnect multiple customer sites with Layer 2 technology.

This topic describes configuring provider edge (PE) switches in an MPLS network using a circuit cross-connect (CCC) on a tagged VLAN interface (802.1Q VLAN) rather than a simple interface.

NOTE: You do not need to make any changes to existing provider switches in your MPLS network to support this type of configuration. For information on configuring provider switches, see [“Configuring MPLS on Provider Switches” on page 88](#).

NOTE: You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

NOTE: If you configure a physical interface as VLAN-tagged and with the `vlan-ccc` encapsulation, you cannot configure the associated logical interfaces with the `inet` family. Doing so could cause the logical interfaces to drop packets.

To configure a PE switch with a VLAN CCC and an MPLS-based Layer 2 circuit:

1. Configure OSPF (or IS-IS) on the loopback (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
user@switch# set ospf area 0.0.0.0 interface interface-name
```

2. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
```



```

user@switch# set interfaces lo0 unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address
user@switch# set interfaces interface-name unit logical-unit-number family inet address address

```

4. Enable the MPLS protocol with CSPF disabled:

NOTE: CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. You need to disable CSPF for link protection to function properly on interarea paths.

```

[edit protocols]
user@switch# set mpls no-cspf

```

5. Configure the customer edge interface as a Layer 2 circuit from the local PE switch to the other PE switch:

```

[edit protocols]
user@switch# set l2circuit neighbor address interface interface-name virtual-circuit-id identifier

```

TIP: Use the switch address of the other switch as the neighbor address.

6. Configure MPLS on the core interfaces:

```

[edit protocols]
user@switch# set mpls interface interface-name
user@switch# set mpls interface interface-name
user@switch# set mpls interface interface-name

```

7. Configure LDP on the loopback interface and the core interfaces:

```

[edit protocols]
user@switch# set ldp interface lo0.0
user@switch# set ldp interface interface-name
user@switch# set ldp interface interface-name
user@switch# set ldp interface interface-name

```

8. Configure **family mpls** on the logical units of the core interfaces:


```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family mpls
user@switch# set interfaces interface-name unit logical-unit-number family mpls
user@switch# set interfaces interface-name unit logical-unit-number family mpls
```

NOTE: You can enable **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

9. Enable VLAN tagging on the customer edge interface of the local PE switch:

```
[edit]
user@switch# set interfaces interface-name vlan-tagging
```

10. Configure the customer edge interface to use VLAN CCC encapsulation:

```
[edit]
user@switch# set interfaces interface-name encapsulation vlan-ccc
```

11. Configure the logical unit of the customer edge interface with a VLAN ID:

NOTE: The VLAN ID cannot be configured on logical interface unit **0**. The logical unit number must be **1** or higher.

The same VLAN ID must be used when configuring the customer edge interface on the other PE switch.

```
[edit ]
user@switch# set interfaces interface-name logical-unit-number vlan-id vlan-id
```

When you have completed configuring one PE switch, follow the same procedures to configure the other PE switch.

NOTE: For EX Series switches, you must use the same type of switch for the other PE switch.

RELATED DOCUMENTATION

[Configuring MPLS on Provider Switches | 88](#)

[Example: Configuring MPLS-Based Layer 2 VPNs | 1438](#)

8

PART

MPLS for Software Defined Network (SDN)

[Introduction to Path Computation Element Protocol \(PCEP\) | 1695](#)

[Configuring PCEP for MPLS RSVP-TE | 1697](#)

[Configuring PCEP for MPLS SPRING-TE | 1777](#)

Introduction to Path Computation Element Protocol (PCEP)

IN THIS CHAPTER

- [PCEP Overview](#) | 1695


PCEP Overview

A Path Computation Element (PCE) is an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. A Path Computation Client (PCC) is any client application requesting a path computation to be performed by a PCE. The Path Computation Element Protocol (PCEP) enables communications between a PCC and a PCE, or between two PCEs (defined in RFC 5440).

PCEP is a TCP-based protocol defined by the IETF PCE Working Group, and defines a set of messages and objects used to manage PCEP sessions and to request and send paths for multidomain traffic engineered LSPs (TE LSPs). It provides a mechanism for a PCE to perform path computation for a PCC's external LSPs. The PCEP interactions include LSP status reports sent by the PCC to the PCE, and PCE updates for the external LSPs.

[Figure 136 on page 1696](#) illustrates the role of PCEP in the client-side implementation of a stateful PCE architecture in an MPLS RSVP-TE enabled network.

Figure 136: PCEP Session



A TCP-based PCEP session connects a PCC to an external PCE. The PCC initiates the PCEP session and stays connected to the PCE for the duration of the PCEP session. During the PCEP session, the PCC requests LSP parameters from the stateful PCE. On receiving one or more LSP parameters from the PCE, the PCC re-signals the TE LSP. When the PCEP session is terminated, the underlying TCP connection is closed immediately, and the PCC attempts to re-establish the PCEP session.

Thus, the PCEP functions include:

- LSP tunnel state synchronization between a PCC and a stateful PCE—When an active stateful PCE connection is detected, a PCC tries to delegate all LSPs to this PCE in a procedure called LSP state synchronization. PCEP enables synchronization of the PCC LSP state to the PCE.
- Delegation of control over LSP tunnels to a stateful PCE—An active stateful PCE controls one or more LSP attributes for computing paths, such as bandwidth, path (ERO), and priority (setup and hold). PCEP enables such delegation of LSPs for path computation.
- Stateful PCE control of timing and sequence of path computations within and across PCEP sessions—An active stateful PCE modifies one or more LSP attributes, such as bandwidth, path (ERO), and priority (setup and hold). PCEP communicates these new LSP attributes from the PCE to the PCC, after which the PCC re-signals the LSP in the specified path.

RELATED DOCUMENTATION

[Support of the Path Computation Element Protocol for RSVP-TE Overview](#) | 1697

Configuring PCEP for MPLS RSVP-TE

IN THIS CHAPTER

- Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697
- Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE | 1717
- Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1734
- Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1746
- Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Controlled Point-to-Multipoint LSPs | 1750
- Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs | 1771

Support of the Path Computation Element Protocol for RSVP-TE Overview

IN THIS SECTION

- Understanding MPLS RSVP-TE | 1698
- Current MPLS RSVP-TE Limitations | 1699
- Use of an External Path Computing Entity | 1701
- Components of External Path Computing | 1701
- Interaction Between a PCE and a PCC Using PCEP | 1703
- LSP Behavior with External Computing | 1707
- Configuration Statements Supported for External Computing | 1709
- PCE-Controlled LSP Protection | 1709
- PCE-Controlled LSP ERO | 1710
- PCE-Controlled Point-to-Multipoint RSVP-TE LSPs | 1710
- PCE-Initiated Point-to-Point LSPs | 1711
- PCE-Initiated Bypass LSP | 1712

- [PCE-Initiated Point-to-Multipoint LSPs | 1713](#)
- [Auto-Bandwidth and PCE-Controlled LSP | 1714](#)
- [TCP-MD5 Authentication for PCEP Sessions | 1714](#)
- [Impact of Client-Side PCE Implementation on Network Performance | 1715](#)

Understanding MPLS RSVP-TE

Traffic engineering (TE) deals with performance optimization of operational networks, mainly mapping traffic flows onto an existing physical topology. Traffic engineering provides the ability to move traffic flow away from the shortest path selected by the interior gateway protocol (IGP) and onto a potentially less congested physical path across a network.

For traffic engineering in large, dense networks, MPLS capabilities can be implemented because they potentially provide most of the functionality available from an overlay model, in an integrated manner, and at a lower cost than the currently competing alternatives. The primary reason for implementing MPLS traffic engineering is to control paths along which traffic flows through a network. The main advantage of implementing MPLS traffic engineering is that it provides a combination of the traffic engineering capabilities of ATM, along with the class-of-service (CoS) differentiation of IP.

In an MPLS network, data plane information is forwarded using label switching. A packet arriving on a provider edge (PE) router from the customer edge (CE) router has labels applied to it, and it is then forwarded to the egress PE router. The labels are removed at the egress router and it is then forwarded out to the appropriate destination as an IP packet. The label-switching routers (LSRs) in the MPLS domain use label distribution protocols to communicate the meaning of labels used to forward traffic between and through the LSRs. RSVP-TE is one such label distribution protocol that enables an LSR peer to learn about the label mappings of other peers.

When both MPLS and RSVP are enabled on a router, MPLS becomes a client of RSVP. The primary purpose of the Junos OS RSVP software is to support dynamic signaling within label-switched paths (LSPs). RSVP reserves resources, such as for IP unicast and multicast flows, and requests quality-of-service (QoS) parameters for applications. The protocol is extended in MPLS traffic engineering to enable RSVP to set up LSPs that can be used for traffic engineering in MPLS networks.

When MPLS and RSVP are combined, labels are associated with RSVP flows. Once an LSP is established, the traffic through the path is defined by the label applied at the ingress node of the LSP. The mapping of label to traffic is accomplished using different criteria. The set of packets that are assigned the same label value by a specific node belong to the same forwarding equivalence class (FEC), and effectively define the RSVP flow. When traffic is mapped onto an LSP in this way, the LSP is called an LSP tunnel.

LSP tunnels are a way to establish unidirectional label-switched paths. RSVP-TE builds on the RSVP core protocol by defining new objects and modifying existing objects used in the PATH and RESV objects for LSP establishment. The new objects—LABEL-REQUEST object (LRO), RECORD-ROUTE object (RRO), LABEL object, and EXPLICIT-ROUTE object (ERO)—are optional with respect to the RSVP protocol, except for the LRO and LABEL objects, which are both mandatory for establishing LSP tunnels.

In general, RSVP-TE establishes a label-switched path that ensures frame delivery from ingress to egress router. However, with the new traffic engineering capabilities, the following functions are supported in an MPLS domain:

- Possibility to establish a label-switched path using either a full or partial explicit route (RFC 3209).
- Constraint-based LSP establishment over links that fulfill requirements, such as bandwidth and link properties.
- Endpoint control, which is associated with establishing and managing LSP tunnels at the ingress and egress routers.
- Link management, which manages link resources to do resource-aware routing of traffic engineering LSPs and to program MPLS labels.
- MPLS fast reroute (FRR), which manages the LSPs that need protection and assigns backup tunnel information to these LSPs.

Current MPLS RSVP-TE Limitations


Although the RSVP extensions for traffic engineering enable better network utilization and meet requirements of classes of traffic, today's MPLS RSVP-TE protocol suite has several issues inherent to its distributed nature. This causes a number of issues during contention for bisection capacity, especially within an LSP priority class where a subset of LSPs share common setup and hold priority values. The limitations of RSVP-TE include:

- Lack of visibility of individual per-LSP, per-device bandwidth demands—The ingress routers in an MPLS RSVP-TE network establish LSPs without having a global view of the bandwidth demand on the network. Information about network resource utilization is only available as total reserved capacity by traffic class on a per interface basis. Individual LSP state is available locally on each label edge router (LER) for its own LSPs only. As a result, a number of issues related to demand pattern arise, particularly within a common setup and hold priority.
- Asynchronous and independent nature of RSVP signaling—In RSVP-TE, the constraints for path establishment are controlled by an administrator. As such, bandwidth reserved for an LSP tunnel is set by the administrator and does not automatically imply any limit on the traffic sent over the tunnel. Therefore, bandwidth available on a traffic engineering link is the bandwidth configured for the link, excluding the sum of all reservations made on the link. Thus, the unsignaled demands on an LSP tunnel lead to service degradation of the LSP requiring excess bandwidth, as well as the other LSPs that comply with the bandwidth requirements of the traffic engineering link.

- LSPs established based on dynamic or explicit path options in the order of preference—The ingress routers in an MPLS RSVP-TE network establish LSPs for demands based on the order of arrival. Because the ingress routers do not have a global view of the bandwidth demand on the network, using the order of preference to establish LSPs can cause traffic to be dropped or LSPs not being established at all when there is an excess of bandwidth demand.

As an example, [Figure 137 on page 1700](#) is configured with MPLS RSVP-TE, in which A and G are the label edge routers (LERs). These ingress routers establish LSPs independently based on the order of demands and have no knowledge or control over each other's LSPs. Routers B, C, and D are intermediate or transit routers that connect to the egress routers E and F.

Figure 137: Example MPLS Traffic Engineering



The ingress routers establish LSPs based on the order in which the demands arrive. If Router G receives two demands of capacity 5 each for G-F, then G signals two LSPs – LSP1 and LSP2 – through G-B-D-F. In the same way, when Router A receives the third demand of capacity 10 for A-E, then it signals an LSP, LSP3, through A-B-C-E. However, if the demand on the A-E LSP increases from 10 to 15, Router A cannot signal LSP3 using the same (A-B-C-E) path, because the B-C link has a lower capacity.

Router A should have signaled the increased demand on LSP3 using the A-B-D-C-E path. Since LSP1 and LSP2 have utilized the B-D link based on the order of demands received, LSP3 is not signaled.

Thus, although adequate max-flow bandwidth is available for all the LSPs, LSP3 is subject to potentially prolonged service degradation. This is due to Router A's lack of global demand visibility and the lack of systemic coordination in demand placement by the ingress routers A and G.

Use of an External Path Computing Entity

As a solution to the current limitations found in the MPLS RSVP-TE path computation, an external path computing entity with a global view of per-LSP, per-device demand in the network independent of available capacity is required.

Currently, only online and real-time constraint-based routing path computation is provided in an MPLS RSVP-TE network. Each router performs constraint-based routing calculations independent of the other routers in the network. These calculations are based on currently available topology information—information that is usually recent, but not completely accurate. LSP placements are locally optimized, based on current network status. The MPLS RSVP-TE tunnels are set up using the CLI. An operator configures the TE LSP, which is then signaled by the ingress router.

In addition to the existing traffic engineering capabilities, the MPLS RSVP-TE functionality is extended to include an external path computing entity, called the Path Computation Element (PCE). The PCE computes the path for the TE LSPs of ingress routers that have been configured for external control. The ingress router that connects to a PCE is called a Path Computation Client (PCC). The PCC is configured with the Path Computation Client Protocol (PCEP) to facilitate external path computing by a PCE.

For more information, see [“Components of External Path Computing” on page 1701](#).

To enable external path computing for a PCC's TE LSPs, include the **lsp-external-controller pccd** statement at the **[edit mpls]** and **[edit mpls lsp lsp-name]** hierarchy levels.

Components of External Path Computing

IN THIS SECTION

- [Path Computation Element | 1701](#)
- [Path Computation Client | 1703](#)
- [Path Computation Element Protocol | 1703](#)

The components that make up an external path computing system are:

Path Computation Element

A Path Computation Element (PCE) can be any entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. However, a PCE can compute the path for only those TE LSPs of a PCC that have been configured for external control.

A PCE can either be stateful or stateless.

- **Stateful PCE**—A stateful PCE maintains strict synchronization between the PCE and network states (in terms of topology and resource information), along with the set of computed paths and reserved resources in use in the network. In other words, a stateful PCE utilizes information from the traffic engineering database as well as information about existing paths (for example, TE LSPs) in the network when processing new requests from the PCC.

A stateful PCE is of two types:

- **Passive stateful PCE**—Maintains synchronization with the PCC and learns the PCC LSP states to better optimize path calculations, but does not have control over them.
- **Active stateful PCE**—Actively modifies the PCC LSPs, in addition to learning about the PCC LSP states.

NOTE: In a redundant configuration with main and backup active stateful PCEs, the backup active stateful PCE cannot modify the attributes of delegated LSPs until it becomes the main PCE at the time of a failover. There is no preempting of PCEs in the case of a switchover. The main PCE is backed by a backup PCE, and when the main PCE goes down, the backup PCE assumes the role of the main PCE and remains the main PCE even after the PCE that was previously the main PCE is operational again.

A stateful PCE provides the following functions:

- Offers offline LSP path computation.
- Triggers LSP re-route when there is a need to re-optimize the network.
- Changes LSP bandwidth when there is an increase in bandwidth demand from an application.
- Modifies other LSP attributes on the router, such as ERO, setup priority, and hold priority.

A PCE has a global view of the bandwidth demand in the network and maintains a traffic-engineered database to perform path computations. It performs statistics collection from all the routers in the MPLS domain using SNMP and NETCONF. This provides a mechanism for offline control of the PCC's TE LSPs. Although an offline LSP path computation system can be embedded in a network controller, the PCE acts like a full-fledged network controller that provides control over the PCC's TE LSPs, in addition to computing paths.

Although a stateful PCE allows for optimal path computation and increased path computation success, it requires reliable state synchronization mechanisms, with potentially significant control plane overhead and the maintenance of a large amount of data in terms of states, as in the case of a full mesh of TE LSPs.

- **Stateless PCE**—A stateless PCE does not remember any computed path, and each set of requests is processed independently of each other (RFC 5440).

Path Computation Client

A Path Computation Client (PCC) is any client application requesting a path computation to be performed by a PCE.

A PCC can connect to a maximum of 10 PCEs at one time. The PCC to PCE connection can be a configured static route or a TCP connection that establishes reachability. The PCC assigns each connected PCE a priority number. It sends a message to all the connected PCEs with information about its current LSPs, in a process called LSP state synchronization. For the TE LSPs that have external control enabled, the PCC delegates those LSPs to the main PCE. The PCC elects, as the main PCE, a PCE with the lowest priority number, or the PCE that it connects to first in the absence of a priority number.

The PCC re-signals an LSP based on the computed path it receives from a PCE. When the PCEP session with the main PCE is terminated, the PCC elects a new main PCE, and all delegated LSPs to the previously main PCE are delegated to the newly available main PCE.

Path Computation Element Protocol

The Path Computation Element Protocol (PCEP) is used for communication between PCC and PCE (as well as between two PCEs) (RFC 5440). PCEP is a TCP-based protocol defined by the IETF PCE Working Group, and defines a set of messages and objects used to manage PCEP sessions and to request and send paths for multidomain TE LSPs. The PCEP interactions include PCC messages, as well as notifications of specific states related to the use of a PCE in the context of MPLS RSVP-TE. When PCEP is used for PCE-to-PCE communication, the requesting PCE assumes the role of a PCC.


Thus, the PCEP functions include:

- LSP tunnel state synchronization between PCC and a stateful PCE.
- Delegation of control over LSP tunnels to a stateful PCE.

Interaction Between a PCE and a PCC Using PCEP

[Figure 138 on page 1704](#) illustrates the relationship between a PCE, PCC, and the role of PCEP in the context of MPLS RSVP-TE.

Figure 138: PCC and RSVP-TE



The PCE to PCC communication is enabled by the TCP-based PCEP. The PCC initiates the PCEP session and stays connected to a PCE for the duration of the PCEP session.

NOTE: Starting with Junos OS Release 16.1, you can secure a PCEP session using TCP-MD5 authentication as per RFC 5440. To enable the MD5 security mechanism for a PCEP session, it is recommended that you define and bind the MD5 authentication key at the **[edit protocols pcep pce pce-id]** hierarchy level for a PCEP session. You can, however, also use a predefined keychain from the **[edit security authentication-key-chains key-chain]** hierarchy level to secure a PCEP session. In this case, you should bind the predefined keychain into the PCEP session at the **[edit protocols pcep pce pce-id]** hierarchy level.

The PCE and PCC use the same key to verify the authenticity of each segment sent on the TCP connection of the PCEP session, thereby securing the PCEP communication between the devices, which might be subject to attacks and can disrupt services on the network.

For more information on securing PCEP sessions using MD5 authentication, see [“TCP-MD5 Authentication for PCEP Sessions” on page 1714](#).

Once the PCEP session is established, the PCC performs the following tasks:

1. LSP state synchronization—The PCC sends information about all the LSPs (local and external) to all connected PCEs. For external LSPs, the PCC sends information about any configuration change, RRO change, state change, and so on, to the PCE.

For PCE-initiated LSPs, there is no LSP configuration present on the PCC. The PCE initiating the LSP sends the LSP parameters to the PCC that has indicated its capability of supporting PCE-initiated LSPs.

NOTE: Support for PCE-initiated LSPs is provided in Junos OS Release 13.3 and later releases.

2. LSP delegation—After the LSP state information is synchronized, the PCC then delegates the external LSPs to one PCE, which is the main active stateful PCE. Only the main PCE can set parameters for the external LSP. The parameters that the main PCE modifies include bandwidth, path (ERO), and priority (setup and hold). The parameters specified in the local configuration are overridden by the parameters that are set by the main PCE.

NOTE: When the PCEP session with the main PCE is terminated, the PCC elects a new main PCE, and all delegated LSPs to the previously main PCE are delegated to the newly available main PCE.

In the case of PCE-initiated LSPs, the PCC creates the LSP using the parameters received from the PCE. The PCC assigns the PCE-initiated LSP a unique LSP-ID, and automatically delegates the LSP to the PCE. A PCC cannot revoke the delegation for the PCE-initiated LSPs for an active PCEP session.

When a PCEP session terminates, the PCC starts two timers without immediately deleting the PCE-initiated LSPs – **delegation cleanup timeout** and **lsp cleanup timer** – to avoid disruption of services. During this time, an active stateful PCE can acquire control of the LSPs provisioned by the failed PCE, by sending a create request for the LSP.

Control over PCE-initiated LSPs reverts to the PCC at the expiration of the **delegation cleanup timeout**. When the **delegation cleanup timeout** expires, and no other PCE has acquired control over the LSP from the failed PCE, the PCC takes local control of the non-delegated PCE-initiated LSP. Later, when the original or a new active stateful PCE wishes to acquire control of the locally controlled PCE-initiated LSPs, the PCC delegates these LSPs to the PCE and the **lsp cleanup timer** timer is stopped.

A PCE may return the delegation of the PCE-initiated LSP to the PCC to allow LSP transfer between PCEs. This triggers the **lsp cleanup timer** for the PCE-initiated LSP. The PCC waits for the LSP cleanup timer to expire before removing the non-delegated PCE-initiated LSPs from the failed PCE.

When the **lsp cleanup timer** expires, and no other PCE has acquired control over the LSPs from the failed PCE, the PCC deletes all the LSPs provisioned by the failed PCE.

NOTE: In compliance with *draft-ietf-pce-stateful-pce-09*, revoking of PCE-initiated LSP delegations by a PCC happens in a make-before-break fashion before the LSPs are redelegated to an alternate PCE. Starting in Junos OS Release 18.1R1, the **lsp-cleanup-timer** must be greater than or equal to the **delegation-cleanup-timeout** for the PCC to revoke the LSP delegations. If not, the redelegation timeout interval for the PCC can be set to infinity, where the LSP delegations to that PCE remain intact until specific action is taken by the PCC to change the parameters set by the PCE.


3. LSP signaling—On receiving one or more LSP parameters from the main active stateful PCE, the PCC re-signals the TE LSP based on the PCE provided path. If the PCC fails to set up the LSP, it notifies the PCE of the setup failure and waits for the main PCE to provide new parameters for that LSP, and then re-signals it.

When the PCE specifies a path that is incomplete or has loose hops where only the path endpoints are specified, the PCC does not perform local constraint-based routing to find out the complete set of hops. Instead, the PCC provides RSVP with the PCE provided path, as it is, for signaling, and the path gets set up using IGP hop-by-hop routing.

Considering the topology used in [Figure 137 on page 1700](#), [Figure 139 on page 1707](#) illustrates the partial client-side PCE implementation in the MPLS RSVP-TE enabled network. The ingress routers A and G are the PCCs that are configured to connect to the external stateful PCE through a TCP connection.

The PCE has a global view of the bandwidth demand in the network and performs external path computations after looking up the traffic engineering database. The active stateful PCE then modifies one or more LSP attributes and sends an update to the PCC. The PCC uses the parameters it receives from the PCE to re-signal the LSP.

Figure 139: Example PCE for MPLS RSVP-TE



This way, the stateful PCE provides a cooperative operation of distributed functionality used to address specific challenges of a shortest interdomain constrained path computation. It eliminates congestion scenarios in which traffic streams are inefficiently mapped onto available resources, causing overutilization of some subsets of network resources, while other resources remain underutilized.

LSP Behavior with External Computing

IN THIS SECTION

- LSP Types | 1707
- LSP Control Mode | 1708

LSP Types

In a client-side PCE implementation, there are three types of TE LSPs:

- CLI-controlled LSPs—The LSPs that do not have the **lsp-external-controller pccd** statement configured are called CLI-controlled LSPs. Although these LSPs are under local control, the PCC updates the connected PCEs with information about the CLI-controlled LSPs during the initial LSP synchronization process. After the initial LSP synchronization, the PCC informs the PCE of any new and deleted LSPs as well.

- PCE-controlled LSPs—The LSPs that have the **lsp-external-controller pccd** statement configured are called PCE-controlled LSPs. The PCC delegates the PCC-initiated LSPs to the main PCE for external path computation.

The PCC informs the PCE about the configured parameters of a PCE-controlled LSP, such as bandwidth, ERO, and priorities. It also informs the PCE about the actual values used for these parameters to set up the LSP including the RRO, when available.

The PCC sends such LSP status reports to the PCE only when a reconfiguration has occurred or when there is a change in the ERO, RRO, or status of the PCE-controlled LSPs under external control.

There are two types of parameters that come from the CLI configuration of an LSP for a PCE:

- Parameters that are not overridden by a PCE, and that are applied immediately.
- Parameters that are overridden by a PCE. These parameters include bandwidth, path, and priority (setup and hold values). When the control mode switches from external to local, the CLI-configured values for these parameters are applied at the next opportunity to re-signal the LSP. The values are not applied immediately.
- Externally-provisioned LSPs (or PCE-initiated LSPs)—The LSPs that have the **lsp-provisioning** statement configured are called PCE-initiated LSPs. A PCE-initiated LSP is dynamically created by an external PCE; as a result, there is no LSP configuration present on the PCC. The PCC creates the PCE-initiated LSP using the parameters provided by the PCE, and automatically delegates the LSP to the PCE.

NOTE: Support for PCE-initiated LSPs is provided in Junos OS Release 13.3 and later releases.

The CLI-controlled LSPs, PCE-controlled LSPs, and PCE-initiated LSPs can coexist on a PCC.

The CLI-controlled LSPs and PCE-controlled LSPs can coexist on a PCC.

LSP Control Mode

In a client-side PCE implementation, there are two types of control modes for a PCC-controlled LSP:

- External—By default, all PCE-controlled LSPs are under external control. When an LSP is under external control, the PCC uses the PCE-provided parameters to set up the LSP.
- Local—A PCE-controlled LSP can come under local control. When the LSP switches from external control to local control, path computation is done using the CLI-configured parameters and constraint-based routing. Such a switchover happens only when there is a trigger to re-signal the LSP. Until then, the PCC uses the PCE-provided parameters to signal the PCE-controlled LSP, although the LSP remains under local control.

A PCE-controlled LSP switches to local control from its default external control mode in cases such as no connectivity to a PCE or when a PCE returns delegation of LSPs back to the PCC.

For more information about CLI-controlled LSPs and PCE-controlled LSPs, see [“LSP Types” on page 1707](#).

Configuration Statements Supported for External Computing

Table 48 on page 1709 lists the MPLS and existing LSP configuration statements that apply to a PCE-controlled LSP.

Table 48: Applicability of MPLS and Existing LSP Configurations to a PCE-Controlled LSP

Support for PCE-Controlled LSP	Applicable LSP Configuration Statements	Applicable MPLS Configuration Statements
These configuration statements can be configured along with the PCE configuration. However, they take effect only when the local configuration is in use. During PCE control, these configuration statements remain inactive.	<ul style="list-style-type: none"> • admin-group • auto-bandwidth • hop-limit • least-fill • most-fill • random 	<ul style="list-style-type: none"> • admin-group • admin-groups • admin-group-extended • hop-limit • no-cspf • smart-optimize-timer
<p>These configuration statements can be configured along with the PCE configuration, but are overridden by the PCE-controlled LSP attributes. However, when the local configuration is in use, the configured values for these configuration statements are applied.</p> <p>NOTE: Changes to the local configuration using the CLI while the LSP is under the control of a stateful PCE do not have any effect on the LSP. These changes come into effect only when the local configuration is applied.</p>	<ul style="list-style-type: none"> • bandwidth • primary • priority 	<ul style="list-style-type: none"> • priority
These configuration statements cannot be configured along with the PCE configuration.	<ul style="list-style-type: none"> • p2mp • template 	<ul style="list-style-type: none"> • p2mp-lsp-next-hop

The rest of the LSP configuration statements are applicable in the same way as for existing LSPs. On configuring any of the above configuration statements for a PCE-controlled LSP, an MPLS log message is generated to indicate when the configured parameters take effect.

PCE-Controlled LSP Protection

The protection paths, including fast reroute and bypass LSPs, are locally computed by the PCC using constraint-based routing. A stateful PCE specifies the primary path (ERO) only. A PCE can also trigger a non-standby secondary path, even if the local configuration does not have a non-standby secondary path for LSP protection.

PCE-Controlled LSP ERO

For PCE-controlled LSPs (PCC-delegated LSPs and PCE-initiated LSPs), only a full-blown Explicit Route Object (ERO) object has to be sent from the PCE to the PCC; otherwise the PCC rejects the PCUpdate or PCCreate message for that PCEP session.

Starting in Junos OS Release 17.2, in addition to **external cspf**, two new path computation types are introduced for the PCE-controlled LSPs: **local cspf** and **no cspf**.

- **local cspf**—A PCC uses the **local cspf** computation type only when the PCE sends in a Juniper Vendor TLV (enterprise number: 0x0a4c) of type 5.
- **no cspf**—Neither the PCE nor the PCC performs a constrained path calculation. The endpoints and constraints are given to the RSVP module for setting up the LSP with the IGP path.

A PCC uses **no cspf** computation type in the following cases:

- When the PCE sends **local cspf** TLV, and when the Junos OS configuration or matching template for this LSP included **no-cspf** in the PCC-delegated LSP.
- When the PCE sends **local cspf** TLV, and when the Junos OS configuration template for this LSP included **no-cspf** in the PCE-initiated LSP.
- When the PCE does not send **local cspf** TLV with an empty ERO or loose ERO (with loose bit set in the ERO object).

With these new computation types, a PCC can accept an ERO object either as a loose ERO, or as an empty ERO. An external path computing entity that is not capable of computing a path can modify parameters such as bandwidth and color, based on the analytics. In such cases, an empty ERO object or loose ERO is used and the path to be taken is decided by the PCC.

PCE-Controlled Point-to-Multipoint RSVP-TE LSPs

After a PCEP session is established between a PCE and a PCC, the PCC reports all the LSPs in the system to the PCE for LSP state synchronization. This includes PCC-controlled, PCE-delegated, and PCE-initiated point-to-point LSPs. Starting with Junos OS Release 15.1F6 and 16.1R1, this capability is extended to report point-to-multipoint LSPs as well. For a PCE, the point-to-multipoint LSP is similar to that of RSVP point-to-multipoint LSP, where the point-to-multipoint LSP is treated as collection of point-to-point LSPs grouped under a point-to-multipoint identifier.

By default, PCE control of point-to-multipoint LSPs is not supported on a PCC. To add this capability, include the **p2mp-lsp-report-capability** statement at the **[edit protocols pcep pce pce-name]** or **[edit protocols pcep pce-group group-id]** hierarchy levels. After the point-to-multipoint report capability is configured on a PCC, the PCC advertises this capability to the PCE. If the PCE advertises the same point-to-multipoint report capability in return, then the PCC reports the complete point-to-multipoint LSP tree to the PCE for LSP state synchronization.

A PCC with the point-to-multipoint TE LSP capability supports reporting of point-to-multipoint TE LSPs for stateful PCEs, point-to-multipoint update, and LSP database supporting point-to-multipoint LSP name as key. However, the following features and functions are not supported for Junos OS Release 15.1F6 and 16.1:

- Static point-to-multipoint LSPs
- PCE-delegated and PCE-initiated point-to-multipoint LSPs
- Auto-bandwidth
- TE++
- PCE request and reply message
- Creation of point-to-multipoint LSPs using templates
- Configuring forward entry on the PCE-initiated point-to-multipoint LSPs
- Configuring forward entry on the router pointing to a provisioned LSP.

PCE-Initiated Point-to-Point LSPs

Starting with Junos OS Release 16.1, the PCEP functionality is extended to allow a stateful PCE to initiate and provision traffic engineering LSPs through a PCC. Earlier, the LSPs were configured on the PCC and the PCC delegated control over the external LSPs to a PCE. The ownership of the LSP state was maintained by the PCC. With the introduction of the PCE-initiated LSPs, a PCE can initiate and provision a traffic engineering point-to-point LSP dynamically without the need for a locally configured LSP on the PCC. On receiving a PCCreate message from a PCE, the PCC creates the PCE-initiated LSP and automatically delegates the LSP to the PCE.

By default, a PCC rejects the request for provisioning PCE-initiated point-to-point LSPs from a PCE. To enable support of PCE-initiated LSPs on the PCC, include the *lsp-provisioning* statement at the **[edit protocols pcep pce pce-id]** or **[edit protocols pcep pce-group group-id]** hierarchy levels.

A PCC indicates its capability of supporting PCE-initiated point-to-point LSPs while establishing the Path Computation Element Protocol (PCEP) session with the PCE. A PCE selects a PCC with this capability to initiate an LSP. The PCE provides the PCC with the PCE-initiated LSP parameters. On receiving the PCE-initiated point-to-point LSP parameters, the PCC sets up the LSP, assigns an LSP ID, and automatically delegates the LSP to the PCE.

When the PCE initiating the LSP does not provide the PCE-initiated point-to-point LSP parameters, the PCC uses the default parameters. An optional LSP template may also be configured to specify values for the PCE-initiated point-to-point LSP when the LSP parameters are not provided by the PCE. To configure an LSP template for PCE-initiated point-to-point LSPs on the PCC, include the [label-switched-path-template](#) statement at the **[edit protocols mpls lsp-external-controller lsp-external-controller]** hierarchy level.

When a PCEP session terminates, the PCC starts two timers without immediately deleting the PCE-initiated LSPs—**delegation cleanup timeout** and **lsp cleanup timer**—to avoid disruption of services. During this time, an active stateful PCE can acquire control of the LSPs provisioned by the failed PCE.

A PCE may return the delegation of the PCE-initiated point-to-point LSP to the PCC to allow LSP transfer between PCEs. Control over PCE-initiated LSPs reverts to the PCC at the expiration of the delegation cleanup timeout. When the delegation cleanup timeout expires, and no other PCE has acquired control over the LSP from the failed PCE, the PCC takes local control of the non-delegated PCE-initiated LSP. Later, when the original or a new active stateful PCE wishes to acquire control of the locally controlled PCE-initiated point-to-point LSPs, the PCC delegates these LSPs to the PCE and the LSP cleanup timer is stopped.

The PCC waits for the LSP cleanup timer to expire before deleting the non-delegated PCE-initiated point-to-point LSPs from the failed PCE. When the LSP cleanup timer expires, and no other PCE has acquired control over the LSPs from the failed PCE, the PCC deletes all the LSPs provisioned by the failed PCE.

PCE-Initiated Bypass LSP

IN THIS SECTION

- [Understanding PCE-Initiated Bypass LSPs | 1712](#)
- [Benefits of PCE-Initiated Bypass LSP | 1713](#)
- [Behavior of PCE-Initiated Bypass LSPs During PCEP Session Failure | 1713](#)

Understanding PCE-Initiated Bypass LSPs

There can be traffic outages at the time of a link or node failure because the backup protection paths in the network do not have sufficient bandwidth to handle traffic. In such networks, although a PCE may be used to compute all the paths, to optimize network performance, the local protection paths also need to be controlled through the PCE.

Junos OS Release 19.2R1 and later releases provide partial support for the Internet draft draft-cbrt-pce-stateful-local-protection-01 (expires December 2018), *PCEP Extensions for RSVP-TE Local-Protection with PCE-Stateful*, where the PCEP functionality is extended to allow a stateful PCE to initiate, provision, and manage bypass LSPs for a protected interface. Multiple bypass LSPs with bandwidth reservation can be initiated by the PCE to protect a link or node. The bandwidth on the bypass LSP is expected to be smaller than the total bandwidth of the primary LSPs that it might protect.

The existing bypass selection mechanism, that prefers manual bypass LSPs (if available) over dynamic bypass LSPs, is extended to prefer PCE-provisioned bypass LSPs (if available) over dynamic bypass LSPs.

The PCE-provisioned bypass LSPs have a higher preference over dynamic bypass LSPs, but are less preferred over manual bypass LSPs.

The set of operations that are used to perform on any operational bypass LSPs, such as **clear rsvp session**, can also be performed on the PCE-initiated bypass LSPs. You can use commands, such as **show path-computation-client status extensive** and **show path-computation-client lsp** to view PCE-initiated bypass LSP statistics.

With the support of PCE-initiated bypass LSP, you can:

- Create a RSVP bypass LSP through PCEP from an external controller, where the bypass LSP:
 - can be for link or node protection.
 - must have a non-zero bandwidth.
 - must have a specified strict ERO.
- Update the bandwidth and ERO for an existing PCE-created bypass LSP.
- Oversubscribe the bypass LSP bandwidth for admission control of primary LSPs. This must be a per-bypass parameter, and should allow updating the subscription per bypass LSP.

Benefits of PCE-Initiated Bypass LSP

The PCE-initiated bypass LSPs provide the following benefits:

- Better control over traffic after a failure and more deterministic path computation of protection paths.
- Meet complex constraints and diversity requirements, such as maintaining diverse paths for LSPs, as well as their local protection paths.
- Ensure links are not overloaded during failure events.

Behavior of PCE-Initiated Bypass LSPs During PCEP Session Failure

At the time of a PCEP session failure, the PCE-initiated bypass LSPs become orphan until the expiration of the state timeout timer. The PCE-initiated bypass LSPs get cleaned up on the expiration of the state timeout timer. To obtain control of a PCE-initiated bypass LSP (after PCEP session fails), a PCE (either the primary PCE, or any secondary PCE) sends a PCInitiate message before the expiration of the state timeout timer.

PCE-Initiated Point-to-Multipoint LSPs

With the introduction of point-to-multipoint PCE-initiated LSPs, a PCE can initiate and provision a point-to-multipoint LSP dynamically without the need for local LSP configuration on the PCC. This enables the PCE to control the timing and sequence of the point-to-multipoint path computations within and across Path Computation Element Protocol (PCEP) sessions, thereby creating a dynamic network that is centrally controlled and deployed.

For more information, see [“Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs” on page 1771.](#)

Auto-Bandwidth and PCE-Controlled LSP

Starting in Junos OS Release 14.2R4, support of auto-bandwidth is provided for PCE-controlled LSPs. In earlier releases, the auto-bandwidth option did not apply to PCE-controlled LSPs, although LSPs under the control of auto-bandwidth and constraint-based routing could coexist with PCE-controlled LSPs. The statistics collection for auto-bandwidth was taking effect only when the control mode of a PCE-controlled LSP changes from external to local. This was happening in cases such as no connectivity to a PCE or when a PCE returns delegation of LSPs back to the PCC.

TCP-MD5 Authentication for PCEP Sessions

A stateful PCE server automates the creation of traffic engineering paths across the network, increasing network utilization and enabling a customized programmable networking experience with the use of PCEP communication with a PCC. A PCC sends LSP reports to a PCE server, and the PCE updates or provisions LSPs back to the PCC. The data sent over a PCEP session is crucial for a PCE server to perform external path computing. As a result, an attack on the PCEP communication can disrupt network services. If altered PCEP messages are sent to a PCC, inappropriate LSPs can be set up. Similarly, if altered PCEP messages are sent to a PCE, an incorrect view of the network is learned by the PCE.

Considering the significance of the PCEP communication between a PCE and PCC in executing the PCE functionalities effectively, Junos OS Release 16.1 introduces the feature of securing a PCEP session using TCP-MD5 authentication as per RFC 5440. This feature protects the communication between a PCE and PCC over a PCEP session, which might be subject to an attack, and can disrupt network services.

To enable the MD5 security mechanism for a PCEP session, it is recommended that you define and bind the MD5 authentication key at the **[edit protocols pcep pce pce-id]** hierarchy level for a PCEP session. You can, however, also use a predefined keychain from the **[edit security authentication-key-chains key-chain]** hierarchy level to secure a PCEP session. In this case, you should bind the predefined keychain into the PCEP session at the **[edit protocols pcep pce pce-id]** hierarchy level.

The following configuration is executed on the PCC to establish a secure PCEP session with a PCE:

- Using MD5 authentication key:

```
[edit protocols pcep pce pce-id]
user@PCC# set authentication-key key
```

- Using predefined authentication keychain:

```
[edit protocols pcep pce pce-id]
```



```

user@PCC# set authentication-key-chain key-chain
user@PCC# set authentication-algorithm md5

```

For secure PCEP sessions to be established successfully, the MD5 authentication should be configured with the pre-shared authentication key on both the PCE server and the PCC. The PCE and PCC use the same key to verify the authenticity of each segment sent on the TCP connection of the PCEP session.

NOTE:

- Junos OS Release 16.1 supports only TCP-MD5 authentication for PCEP sessions, without extending support for TLS and TCP-AO, such as protection against eavesdropping, tampering, and message forgery.
- Initial application of security mechanism to a PCEP session causes the session to reset.
- If MD5 is misconfigured or not configured on one side of the PCEP session, the session does not get established. Verify that the configurations on the PCC and PCE are matching.
- This feature does not provide support for any session authentication mechanism.
- To view the authentication keychain used by the PCEP session, use the **show path-computation-client status** and **show protocols pcep** command outputs.
- Use the **show system statistics tcp | match auth** command to view the number of packets that get dropped by TCP because of authentication errors.
- Operation of the keychain can be verified by using the **show security keychain detail** command output.

Impact of Client-Side PCE Implementation on Network Performance

The maintenance of a stateful database can be non-trivial. In a single centralized PCE environment, a stateful PCE simply needs to remember all the TE LSPs that the PCE has computed, the TE LSPs that were actually set up (if this can be known), and when the TE LSPs were torn down. However, these requirements cause substantial control protocol overhead in terms of state, network usage and processing, and optimizing links globally across the network. Thus, the concerns of a stateful PCE implementation include:

- Any reliable synchronization mechanism results in significant control plane overhead. PCEs might synchronize state by communicating with each other, but when TE LSPs are set up using distributed computation performed among several PCEs, the problems of synchronization and race condition avoidance become larger and more complex.
- Out-of-band traffic engineering database synchronization can be complex with multiple PCEs set up in a distributed PCE computation model, and can be prone to race conditions, scalability concerns, and so on.

- Path calculations incorporating total network state is highly complex, even if the PCE has detailed information on all paths, priorities, and layers.

In spite of the above concerns, the partial client-side implementation of the stateful PCE is extremely effective in large traffic engineering systems. It provides rapid convergence and significant benefits in terms of optimal resource usage, by providing the requirement for global visibility of a TE LSP state and for ordered control of path reservations across devices within the system being controlled.

Release History Table

Release	Description
17.2R1	Starting in Junos OS Release 17.2, in addition to external cspf , two new path computation types are introduced for the PCE-controlled LSPs: local cspf and no cspf .
16.1	Starting with Junos OS Release 16.1, you can secure a PCEP session using TCP-MD5 authentication as per RFC 5440.
16.1	Junos OS Release 16.1 introduces the feature of securing a PCEP session using TCP-MD5 authentication as per RFC 5440.
14.2R4	Starting in Junos OS Release 14.2R4, support of auto-bandwidth is provided for PCE-controlled LSPs.

RELATED DOCUMENTATION

[Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE | 1717](#)

[Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1734](#)

[Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Controlled Point-to-Multipoint LSPs | 1750](#)

Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE

IN THIS SECTION

- [Requirements | 1717](#)
- [Overview | 1718](#)
- [Configuration | 1720](#)
- [Verification | 1727](#)

This example shows how to enable external path computing by a Path Computation Element (PCE) for traffic engineered label-switched paths (TE LSPs) on a Path Computation Client (PCC). It also shows how to configure the Path Computation Element Protocol (PCEP) on the PCC to enable PCE to PCC communication.

Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of ACX Series routers, M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, T Series Core Routers, or PTX Series Transport Router, one of which is configured as a PCC.
- A TCP connection to an external stateful PCE from the PCC.
- Junos OS Release 12.3 or later running on the PCC along with the JSDN add-on package.

NOTE: The JSDN add-on package is required to be installed along with the core Junos OS installation package.

Before you begin:

1. Configure the device interfaces.
2. Configure MPLS and RSVP-TE.
3. Configure IS-IS or any other IGP protocol.

Overview

Starting with Junos OS Release 12.3, the MPLS RSVP-TE functionality is extended to provide a partial client-side implementation of the stateful PCE architecture (draft-ietf-pce-stateful-pce) on a PCC.

NOTE: The partial client-side implementation of the stateful PCE architecture is based on version 2 of Internet draft draft-ietf-pce-stateful-pce. Starting with Junos OS Release 16.1, this implementation is upgraded to support version 7, as defined in Internet draft draft-ietf-pce-stateful-pce-07. Releases prior to 16.1 support the older version of the PCE draft, causing interoperability issues between a PCC running a previous release and a stateful PCE server that adheres to Internet draft draft-ietf-pce-stateful-pce-07.

To enable external path computing by a PCE, include the **lsp-external-controller** statement on the PCC at the **[edit mpls]** and **[edit mpls lsp *lsp-name*]** hierarchy levels.

```
lsp-external-controller pccd;
```

An LSP configured with the **lsp-external-controller** statement is referred to as a PCE-controlled LSP and is under the external control of a PCE by default. An active stateful PCE can override the parameters set from the CLI, such as bandwidth, path (ERO), and priority, for such PCE-controlled LSPs of the PCC.

To enable PCE to PCC communication, configure PCEP on the PCC at the **[edit protocols]** hierarchy level.

```
pcep { ... }
```

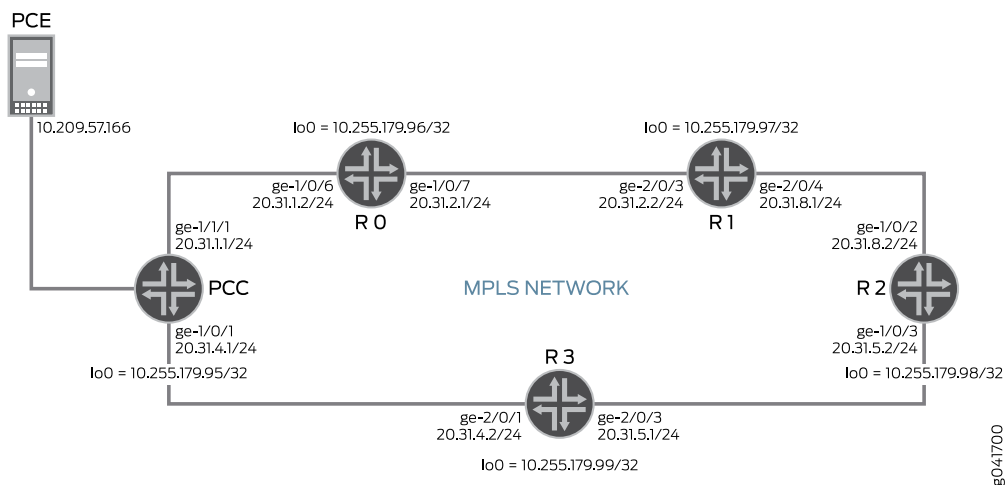
When configuring PCEP on a PCC, be aware of the following considerations:

- The JSDN add-on package is required to be installed along with the core Junos OS installation package.
- Junos OS Release 12.3 supports only stateful PCEs.
- A PCC can connect to a maximum of 10 stateful PCEs. At any given point in time, there can be only one main PCE (the PCE with the lowest priority value, or the PCE that connects to the PCC first in the absence of a PCE priority) to which the PCC delegates LSPs for path computation.
- For Junos OS Release 12.3, the PCC always initiates the PCEP sessions. PCEP sessions initiated by remote PCEs are not accepted by the PCC.
- Existing LSP features, such as LSP protection and make-before-break, work for PCE-controlled LSPs.
- The auto-bandwidth option is turned off for PCE-controlled LSPs, although LSPs under the control of auto-bandwidth and constraint-based routing can coexist with PCE-controlled LSPs.
- PCE-controlled LSPs can be referred to by other CLI configurations, such as `lsp-nexthop` to routes, forwarding adjacencies, CCC connections, and logical tunnels.

- PCE-controlled LSPs do not support GRES.
- PCE-controlled LSPs under logical-systems are not supported.
- PCE-controlled LSPs cannot be point-to-multipoint LSPs.
- Bidirectional LSPs are not supported.
- PCE-controlled LSPs cannot have secondary paths without a primary path.
- PCE-controlled LSPs depend on external path computation, which impacts overall setup time, reroutes, and make-before-break features.
- Setup time and convergence time (reroute, MBB) for existing LSPs is the same as in previous releases, in the absence of PCE-controlled LSPs. However, a small impact is seen in the presence of PCE-controlled LSPs.
- ERO computation time is expected to be significantly higher than local-CSPF.

Topology

Figure 140: Configuring PCEP for MPLS RSVP-TE



In this example, PCC is the ingress router that connects to the external active stateful PCE.

The external LSPs of Router PCC are computed as follows:

1. Router PCC receives the LSP tunnel configuration that was set up using the CLI. Assuming that the received configuration is enabled with external path computing, Router PCC becomes aware that some of the LSP attributes – bandwidth, path, and priority – are under the control of the stateful PCE and delegates the LSP to the PCE.

In this example, the external LSP is called **PCC-to-R2** and it is being set up from Router PCC to Router R2. The CLI-configured ERO for **PCC-to-R2** is PCC-R0-R1-R2. The bandwidth for **PCC-to-R2** is 10m, and both the setup and hold priority values are 4.

2. Router PCC tries to retrieve the PCE-controlled LSP attributes. To do this, Router PCC sends out a PCRpt message to the stateful PCE stating that the LSP has been configured. The PCRpt message communicates the status of the LSP and contains the local configuration parameters of the LSP.
3. The stateful PCE modifies one or more of the delegated LSP attributes and sends the new LSP parameters to Router PCC through the PCUpd message.
4. On receiving the new LSP parameters, Router PCC sets up a new LSP and re-signals it using the PCE-provided path.

In this example, the PCE-provided ERO for **PCC-to-R2** is PCC-R3-R2. The bandwidth for **PCC-to-R2** is 8m, and both the setup and hold priority values are 3.

5. Router PCC sends a PCRpt with the new RRO to the stateful PCE.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

PCC

```
set interfaces ge-1/0/1 unit 0 family inet address 20.31.4.1/24
set interfaces ge-1/0/1 unit 0 family iso
set interfaces ge-1/0/1 unit 0 family mpls
set interfaces ge-1/1/1 unit 0 family inet address 20.31.1.1/24
set interfaces ge-1/1/1 unit 0 family iso
set interfaces ge-1/1/1 unit 0 family mpls
```



```

set interfaces lo0 unit 0 family inet address 10.255.179.95/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls lsp-external-controller pccd
set protocols mpls label-switched-path PCC-to-R2 to 10.255.179.98
set protocols mpls label-switched-path PCC-to-R2 bandwidth 10m
set protocols mpls label-switched-path PCC-to-R2 priority 4 4
set protocols mpls label-switched-path PCC-to-R2 primary to-R2-path
set protocols mpls label-switched-path PCC-to-R2 lsp-external-controller pccd
set protocols mpls path to-R2-path 20.31.1.2 strict
set protocols mpls path to-R2-path 20.31.2.2 strict
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis level 1 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0
set protocols pcep pce pce1 destination-ipv4-address 10.209.57.166
set protocols pcep pce pce1 destination-port 4189
set protocols pcep pce pce1 pce-type active
set protocols pcep pce pce1 pce-type stateful

```

R0

```

set interfaces ge-1/0/6 unit 0 family inet address 20.31.1.2/24
set interfaces ge-1/0/6 unit 0 family iso
set interfaces ge-1/0/6 unit 0 family mpls
set interfaces ge-1/0/7 unit 0 family inet address 20.31.2.1/24
set interfaces ge-1/0/7 unit 0 family iso
set interfaces ge-1/0/7 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.179.96/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis level 1 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0

```


R1

```
set system ports console log-out-on-disconnect
set interfaces ge-2/0/3 unit 0 family inet address 20.31.2.2/24
set interfaces ge-2/0/3 unit 0 family iso
set interfaces ge-2/0/3 unit 0 family mpls
set interfaces ge-2/0/4 unit 0 family inet address 20.31.8.1/24
set interfaces ge-2/0/4 unit 0 family iso
set interfaces ge-2/0/4 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.179.97/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis level 1 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0
```

R2

```
set interfaces ge-1/0/2 unit 0 family inet address 20.31.8.2/24
set interfaces ge-1/0/2 unit 0 family iso
set interfaces ge-1/0/2 unit 0 family mpls
set interfaces ge-1/0/3 unit 0 family inet address 20.31.5.2/24
set interfaces ge-1/0/3 unit 0 family iso
set interfaces ge-1/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.179.98/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis level 1 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0
```

R3


```

set interfaces ge-2/0/1 unit 0 family inet address 20.31.4.2/24
set interfaces ge-2/0/1 unit 0 family iso
set interfaces ge-2/0/1 unit 0 family mpls
set interfaces ge-2/0/3 unit 0 family inet address 20.31.5.1/24
set interfaces ge-2/0/3 unit 0 family iso
set interfaces ge-2/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.179.99/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis level 1 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure Router PCC:

NOTE: Repeat this procedure for every Juniper Networks ingress router in the MPLS domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure the interfaces.

To enable MPLS, include the protocol family on the interface so that the interface does not discard incoming MPLS traffic.

```

[edit interfaces]
user@PCC# set ge-1/0/1 unit 0 family inet address 20.31.4.1/24
user@PCC# set ge-1/0/1 unit 0 family iso
user@PCC# set ge-1/0/1 unit 0 family mpls
user@PCC# set ge-1/1/1 unit 0 family inet address 20.31.1.1/24
user@PCC# set ge-1/1/1 unit 0 family iso
user@PCC# set ge-1/1/1 unit 0 family mpls
user@PCC# set lo0 unit 0 family inet address 10.255.179.95/32

```


2. Enable RSVP on all interfaces of Router PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set rsvp interface all
user@PCC# set rsvp interface fxp0.0 disable
```

3. Configure the label-switched path (LSP) from Router PCC to Router R2 and enable external control of LSPs by the PCE.

```
[edit protocols]
user@PCC# set mpls lsp-external-controller pccd
user@PCC# set mpls label-switched-path PCC-to-R2 to 10.255.179.98/32
user@PCC# set mpls label-switched-path PCC-to-R2 bandwidth 10m
user@PCC# set protocols mpls label-switched-path PCC-to-R2 priority 4 4
user@PCC# set protocols mpls label-switched-path PCC-to-R2 primary to-R2-path
user@PCC# set protocols mpls label-switched-path PCC-to-R2 lsp-external-controller pccd
```

4. Configure the LSP from Router PCC to Router R2, which has local control and is overridden by the PCE-provided LSP parameters.

```
[edit protocols]
user@PCC# set mpls path to-R2-path 20.31.1.2/30 strict
user@PCC# set mpls path to-R2-path 20.31.2.2/30 strict
```

5. Enable MPLS on all interfaces of Router PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set mpls interface all
user@PCC# set mpls interface fxp0.0 disable
```

6. Configure IS-IS on all interfaces of Router PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set isis level 1 disable
user@PCC# set isis interface all
user@PCC# set isis interface fxp0.0 disable
user@PCC# set isis interface lo0.0
```

7. Define the PCE that Router PCC connects to, and configure the IP address of the PCE.


```
[edit protocols]
user@PCC# set pcep pce pce1 destination-ipv4-address 10.209.57.166
```

8. Configure the destination port for Router PCC that connects to a PCE using the TCP-based PCEP.

```
[edit protocols]
user@PCC# set pcep pce pce1 destination-port 4189
```

9. Configure the PCE type.

```
[edit protocols]
user@PCC# set pcep pce pce1 pce-type active
user@PCC# set pcep pce pce1 pce-type stateful
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PCC# show interfaces
ge-1/0/1 {
  unit 0 {
    family inet {
      address 20.31.4.1/24;
    }
    family iso;
    family mpls;
  }
}
ge-1/1/1 {
  unit 0 {
    family inet {
      address 20.31.1.1/24;
    }
    family iso;
    family mpls;
  }
}
lo0 {
```



```

unit 0 {
    family inet {
        address 10.255.179.95/32;
    }
}

```

user@PCC# **show protocols**

```

rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    lsp-external-controller pccd;
    label-switched-path PCC-to-R2 {
        to 10.255.179.98;
        bandwidth 10m;
        priority 4 4;
        primary to-R2-path;
        lsp-external-controller pccd;
    }
    path to-R2-path {
        20.31.1.2 strict;
        20.31.2.2 strict;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
}
isis {
    level 1 disable;
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0;
}
pcep {
    pce pce1 {
        destination-ipv4-address 10.209.57.166;
        destination-port 4189;
    }
}

```



```

    pce-type active stateful;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying the PCEP Session Status | 1727](#)
- [Verifying the PCE-Controlled LSP Status When LSP Control Is External | 1728](#)
- [Verifying the PCE-Controlled LSP Status When LSP Control Is Local | 1730](#)

Confirm that the configuration is working properly.

Verifying the PCEP Session Status

Purpose

Verify the PCEP session status between the PCE and Router PCC when the PCE status is up.

Action

From operational mode, run the **show path-computation-client active-pce** command.

user@PCC> show path-computation-client active-pce

```

PCE pce1
General
  IP address           : 10.209.57.166
  Priority              : 0
  PCE status           : PCE_STATE_UP
  Session type         : PCE_TYPE_STATEFULACTIVE
  PCE-mastership       : main

Counters
  PCReqs               Total: 0           last 5min: 0           last hour: 0
  PCReps               Total: 0           last 5min: 0           last hour: 0

```


PCRpts	Total: 5	last 5min: 5	last hour: 5
PCUpdates	Total: 1	last 5min: 1	last hour: 1
Timers			
Local	Keepalive timer:	30 [s]	Dead timer: 120 [s]
Remote	Keepalive timer:	30 [s]	Dead timer: 120 [s]
Errors			
PCErr-recv			
PCErr-sent			
PCE-PCC-NTFS			
PCC-PCE-NTFS			

Meaning

The output displays information about the current active stateful PCE that Router PCC is connected to. The **PCE status** output field indicates the current status of the PCEP session between the PCE and Router PCC.

For **pce1**, the status of the PCEP session is **PCE_STATE_UP**, which indicates that the PCEP session has been established between the PCEP peers.

The statistics of **PCRpts** indicate the number of messages sent by Router PCC to the PCE to report the current status of LSPs. The **PCUpdates** statistics indicate the number of messages received by Router PCC from the PCE. The **PCUpdates** messages include the PCE modified parameters for the PCE-controlled LSPs.

Verifying the PCE-Controlled LSP Status When LSP Control Is External

Purpose

Verify the status of the PCE-controlled LSP from Router PCC to Router R2 when the LSP is under external control.

Action

From operational mode, run the **show mpls lsp name PCC-to-R2 extensive** command.

user@PCC> show mpls lsp name PCC-to-R2 extensive

```
Ingress LSP: 1 sessions

10.255.179.98
  From: 10.255.183.59, State: Up, ActiveRoute: 0, LSPname: PCC-to-R2
```



```

ActivePath: to-R2-path (primary)
LSPtype: Externally controlled, Penultimate hop popping
LSP Control Status: Externally controlled
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary   to-R2-path      State: Up
  Priorities: 3 3
  Bandwidth: 8Mbps
  SmartOptimizeTimer: 180
    No computed ERO.
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    20.31.4.2 20.31.5.2
21 Mar 11 05:00:56.736 EXTCTRL LSP: Sent Path computation request and LSP status

20 Mar 11 05:00:56.736 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
19 Mar 11 05:00:56.735 Selected as active path
18 Mar 11 05:00:56.734 EXTCTRL LSP: Sent Path computation request and LSP status

17 Mar 11 05:00:56.734 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
16 Mar 11 05:00:56.734 Record Route: 20.31.4.2 20.31.5.2
15 Mar 11 05:00:56.734 Up
14 Mar 11 05:00:56.713 EXTCTRL LSP: Sent Path computation request and LSP status

13 Mar 11 05:00:56.713 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
12 Mar 11 05:00:56.712 Originate Call
11 Mar 11 05:00:56.712 EXTCTRL_LSP: Received setup parameters : 20.31.4.2
20.31.5.2
10 Mar 11 05:00:49.283 EXTCTRL LSP: Sent Path computation request and LSP status

9 Mar 11 05:00:49.283 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
8 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
7 Mar 11 05:00:20.581 EXTCTRL LSP: Sent Path computation request and LSP status

6 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
5 Mar 11 05:00:20.580 EXTCTRL_LSP: Control status became external
4 Mar 11 05:00:03.716 EXTCTRL_LSP: Control status became local
3 Mar 11 05:00:03.714 EXTCTRL LSP: Sent Path computation request and LSP status

```



```

2 Mar 11 05:00:03.714 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
1 Mar 11 05:00:00.279 EXTCTRL LSP: Awaiting external controller connection
Created: Mon Mar 11 05:00:00 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

In the output, the **LSPTYPE** and **LSP Control Status** output fields show that the LSP is externally controlled. The output also shows a log of the PCEP messages sent between Router PCC and the PCE.

The PCEP session between the PCE and Router PCC is up, and Router PCC receives the following PCE-controlled LSP parameters:

- ERO (path)—20.31.4.2 and 20.31.5.2
- Bandwidth—8Mbps
- Priorities—3 3 (setup and hold values)

Verifying the PCE-Controlled LSP Status When LSP Control Is Local

Purpose

Verify the status of the PCE-controlled LSP from Router PCC to Router R2 when the LSP control becomes local.

Action

From operational mode, run the **show mpls lsp name PCC-to-R2 extensive** command.

user@PCC> show mpls lsp name PCC-to-R2 extensive

```

Ingress LSP: 1 sessions

10.255.179.98
From: 10.255.183.59, State: Up, ActiveRoute: 0, LSPname: PCC-to-R2
ActivePath: to-R2-path (primary)
LSPTYPE: Externally controlled, Penultimate hop popping
LSP Control Status: Locally controlled

```



```

LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary to-R2-path State: Up
Priorities: 4 4 (ActualPriorities 3 3)
Bandwidth: 10Mbps (ActualBandwidth: 8Mbps)
SmartOptimizeTimer: 180
No computed ERO.
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    20.31.4.2 20.31.5.2
22 Mar 11 05:02:09.618 EXTCTRL_LSP: Control status became local
21 Mar 11 05:00:56.736 EXTCTRL LSP: Sent Path computation request and LSP status

20 Mar 11 05:00:56.736 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
19 Mar 11 05:00:56.735 Selected as active path
18 Mar 11 05:00:56.734 EXTCTRL LSP: Sent Path computation request and LSP status

17 Mar 11 05:00:56.734 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
16 Mar 11 05:00:56.734 Record Route: 20.31.4.2 20.31.5.2
15 Mar 11 05:00:56.734 Up
14 Mar 11 05:00:56.713 EXTCTRL LSP: Sent Path computation request and LSP status

13 Mar 11 05:00:56.713 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
12 Mar 11 05:00:56.712 Originate Call
11 Mar 11 05:00:56.712 EXTCTRL_LSP: Received setup parameters : 20.31.4.2
20.31.5.2
10 Mar 11 05:00:49.283 EXTCTRL LSP: Sent Path computation request and LSP status

9 Mar 11 05:00:49.283 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
8 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
7 Mar 11 05:00:20.581 EXTCTRL LSP: Sent Path computation request and LSP status

6 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
5 Mar 11 05:00:20.580 EXTCTRL_LSP: Control status became external
4 Mar 11 05:00:03.716 EXTCTRL_LSP: Control status became local
3 Mar 11 05:00:03.714 EXTCTRL LSP: Sent Path computation request and LSP status

2 Mar 11 05:00:03.714 EXTCTRL_LSP: Computation request/lsp status contains:

```



```
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
  1 Mar 11 05:00:00.279 EXTCTRL LSP: Awaiting external controller connection
    Created: Mon Mar 11 05:00:00 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

In the output, the **LSP Control Status** output field shows that the LSP is under local control. Although the PCE-controlled LSP is under local control, Router PCC continues to use the PCE-provided parameters, until the next opportunity to re-signal the LSP.

The output now displays the LSP parameters that were configured using the CLI along with the PCE-provided parameters used to establish the LSP as the actual values in use.

- Bandwidth—10Mbps (ActualBandwidth: 8Mbps)
- Priorities—4 4 (ActualPriorities 3 3)

On the trigger to re-signal the LSP, Router PCC uses the local configuration parameters to establish the PCE-controlled LSP.

user@PCC> show mpls lsp name PCC-to-R2 extensive externally-controlled

```
Ingress LSP: 1 sessions

10.255.179.98
  From: 10.255.183.59, State: Up, ActiveRoute: 0, LSPname: PCC-to-R2
  ActivePath: to-R2-path (primary)
  LSPTYPE: Externally controlled, Penultimate hop popping
  LSP Control Status: Locally controlled
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary to-R2-path State: Up
  Priorities: 4 4
  Bandwidth: 10Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
20.31.1.2 S 20.31.2.2 S 20.31.8.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
```



```

20=Node-ID):
    20.31.1.2 20.31.2.2 20.31.8.2
28 Mar 11 05:02:51.787 Record Route: 20.31.1.2 20.31.2.2 20.31.8.2
27 Mar 11 05:02:51.787 Up
26 Mar 11 05:02:51.697 EXTCTRL_LSP: Applying local parameters with this
signalling attempt
25 Mar 11 05:02:51.697 Originate Call
24 Mar 11 05:02:51.696 Clear Call
23 Mar 11 05:02:51.696 CSPF: computation result accepted 20.31.1.2 20.31.2.2
20.31.8.2
22 Mar 11 05:02:09.618 EXTCTRL_LSP: Control status became local
21 Mar 11 05:00:56.736 EXTCTRL LSP: Sent Path computation request and LSP status

20 Mar 11 05:00:56.736 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
19 Mar 11 05:00:56.735 Selected as active path
18 Mar 11 05:00:56.734 EXTCTRL LSP: Sent Path computation request and LSP status

17 Mar 11 05:00:56.734 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
16 Mar 11 05:00:56.734 Record Route: 20.31.4.2 20.31.5.2
15 Mar 11 05:00:56.734 Up
14 Mar 11 05:00:56.713 EXTCTRL LSP: Sent Path computation request and LSP status

13 Mar 11 05:00:56.713 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
12 Mar 11 05:00:56.712 Originate Call
11 Mar 11 05:00:56.712 EXTCTRL_LSP: Received setup parameters : 20.31.4.2
20.31.5.2
10 Mar 11 05:00:49.283 EXTCTRL LSP: Sent Path computation request and LSP status

9 Mar 11 05:00:49.283 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
8 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
7 Mar 11 05:00:20.581 EXTCTRL LSP: Sent Path computation request and LSP status

6 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
5 Mar 11 05:00:20.580 EXTCTRL_LSP: Control status became external
4 Mar 11 05:00:03.716 EXTCTRL_LSP: Control status became local
3 Mar 11 05:00:03.714 EXTCTRL LSP: Sent Path computation request and LSP status

2 Mar 11 05:00:03.714 EXTCTRL_LSP: Computation request/lsp status contains:

```



```
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
  1 Mar 11 05:00:00.279 EXTCTRL LSP: Awaiting external controller connection
    Created: Mon Mar 11 05:00:00 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

The **Computed ERO** is 20.31.1.2, 20.31.2.2, and 20.31.8.2. The PCE-controlled LSP is established using the local configuration parameters.

RELATED DOCUMENTATION

[Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697](#)

Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs

IN THIS SECTION

- [Requirements | 1735](#)
- [Overview | 1735](#)
- [Configuration | 1737](#)
- [Verification | 1742](#)

This example shows how to configure the Path Computation Client (PCC) with the capability of supporting Path Computation Element (PCE)-initiated traffic-engineered point-to-point label-switched paths (LSPs).

Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of ACX Series, M Series, MX Series, or T Series routers.
- A TCP connection to two external stateful PCEs from the ingress router (PCC).
- Junos OS Release 16.1 or later running on the PCC.

Before you begin:

- Configure the device interfaces.
- Configure MPLS and RSVP-TE (RSVP-Traffic Engineering).
- Configure OSPF or any other IGP protocol.

Overview

Starting with Junos OS Release 16.1, the PCEP functionality is extended to allow a stateful PCE to initiate and provision traffic engineering LSPs through a PCC. Earlier, the LSPs were configured on the PCC and the PCC delegated control over the external LSPs to a PCE. The ownership of the LSP state was maintained by the PCC. With the introduction of the PCE-initiated LSPs, a PCE can initiate and provision a traffic engineering point-to-point LSP dynamically without the need for a locally configured LSP on the PCC. On receiving a PCCreate message from a PCE, the PCC creates the PCE-initiated LSP and automatically delegates the LSP to the PCE.

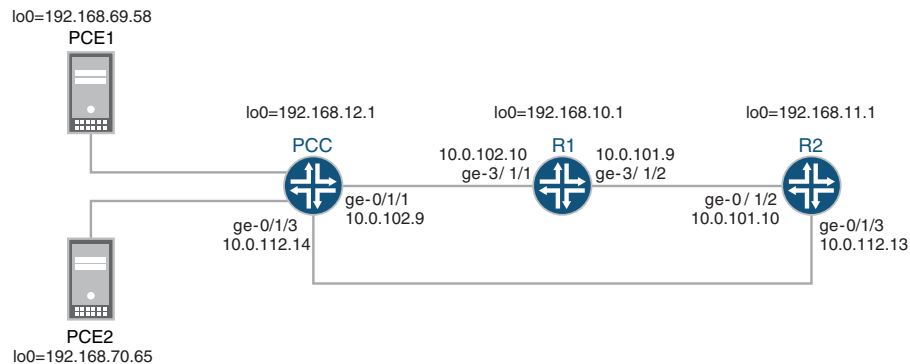
When configuring the support of PCE-initiated point-to-point LSPs for a PCC, be aware of the following considerations:

- Junos OS Release 13.3 supports only stateful PCEs.
- For Junos OS Release 13.3, the PCC always initiates the PCEP sessions. PCEP sessions initiated by remote PCEs are not accepted by the PCC.
- Existing LSP features, such as LSP protection and make-before-break, work for PCE-initiated LSPs.
- PCE-initiated LSPs do not support graceful Routing Engine switchover (GRES).
- PCE-initiated LSPs under logical systems are not supported.
- PCE-initiated LSPs cannot be point-to-multipoint LSPs.
- Bidirectional LSPs are not supported.
- RSVP-TE for unnumbered links is not supported. PCE-initiated LSPs support only numbered links.
- The PCE initiating a segment routing LSP can use the binding segment ID (SID) labels associated with non-colored segment routing LSPs to provision the PCE-initiated segment routing LSP paths.

Starting in Junos OS Release 18.2R1, statically configured non-colored segment routing LSPs on the ingress device are reported to a PCE through a PCEP session. These non-colored segment routing LSPs may have binding SID labels associated with them. With this feature, the PCE can use this binding SID label in the label stack to provision PCE-initiated segment routing LSP paths.

Topology

Figure 141: Example PCE-Initiated Point-to-Point LSP for MPLS RSVP-TE



In this example, PCC is the ingress router that connects to two external stateful PCEs: PCE1 and PCE2.

When there is a new demand, the active stateful PCE dynamically initiates an LSP to meet the requirement. Since PCC is configured with the capability of supporting the PCE-initiated LSP, path computation on PCC is performed as follows:

1. A PCE sends a PCCreate message to the PCC to initiate and provision an LSP. The PCC sets up the PCE-initiated LSP using the parameters received from the PCE, and automatically delegates the PCE-initiated LSP to the PCE that initiated it.

In this example, PCE1 is the active stateful PCE that initiates and provisions the PCE-initiated LSP on PCC. On receiving the PCE-initiated LSP parameters, PCC sets up the LSP and automatically delegates the PCE-initiated LSP to PCE1.

2. When the PCEP session between PCC and PCE1 is terminated, PCC starts two timers for the PCE1-initiated LSP: delegation cleanup timeout and the LSP cleanup timer. During this time, PCE1 or PCE2 can acquire control of the PCE-initiated LSP.
3. If PCE2 acquires control over the PCE-initiated LSP before the expiration of the LSP cleanup timer, PCC delegates the PCE-initiated LSP to PCE2, and the LSP cleanup timer and the delegation cleanup timeout are stopped.

4. If the delegation cleanup timeout expired, and neither PCE1 nor PCE2 acquired control over the PCE-initiated LSP, PCC takes local control of the non-delegated PCE-initiated LSP until the expiration of the LSP cleanup timer.
5. After the expiration of the LSP cleanup timer, PCC deletes the PCE-initiated LSP provisioned by PCE1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

PCC

```

set interfaces ge-0/1/1 unit 0 family inet address 10.0.102.9/24
set interfaces ge-0/1/1 unit 0 family iso
set interfaces ge-0/1/1 unit 0 family mpls
set interfaces ge-0/1/3 unit 0 family inet address 10.0.112.14/24
set interfaces ge-0/1/3 unit 0 family iso
set interfaces ge-0/1/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.12.1/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls lsp-external-controller ppcd
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols pcep pce-group PCEGROUP pce-type active
set protocols pcep pce-group PCEGROUP pce-type stateful
set protocols pcep pce-group PCEGROUP lsp-provisioning
set protocols pcep pce-group PCEGROUP lsp-cleanup-timer 30
set protocols pcep pce PCE1 destination-ipv4-address 192.168.69.58
set protocols pcep pce PCE1 destination-port 4189
set protocols pcep pce PCE1 pce-group PCEGROUP
set protocols pcep pce PCE2 destination-ipv4-address 192.168.70.65
set protocols pcep pce PCE2 destination-port 4189
set protocols pcep pce PCE2 pce-group PCEGROUP

```

R1


```

set interfaces ge-3/1/1 unit 0 family inet address 10.0.102.10/24
set interfaces ge-3/1/1 unit 0 family iso
set interfaces ge-3/1/1 unit 0 family mpls
set interfaces ge-3/1/2 unit 0 family inet address 10.0.101.9/24
set interfaces ge-3/1/2 unit 0 family iso
set interfaces ge-3/1/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.10.1/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

R2

```

set interfaces ge-0/1/1 unit 0 family inet address 10.0.101.10/24
set interfaces ge-0/1/1 unit 0 family iso
set interfaces ge-0/1/1 unit 0 family mpls
set interfaces ge-0/1/3 unit 0 family inet address 10.0.112.13/24
set interfaces ge-0/1/3 unit 0 family iso
set interfaces ge-0/1/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.11.1/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the PCC router:

NOTE: Repeat this procedure for every Juniper Networks ingress router in the MPLS domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure the interfaces.

To enable MPLS, include the protocol family on the interface so that the interface does not discard incoming MPLS traffic.

```
[edit interfaces]
user@PCC# set ge-0/1/1 unit 0 family inet address 10.0.102.9/24
user@PCC# set ge-0/1/1 unit 0 family iso
user@PCC# set ge-0/1/1 unit 0 family mpls
user@PCC# set ge-0/1/3 unit 0 family inet address 10.0.112.14/24
user@PCC# set ge-0/1/3 unit 0 family iso
user@PCC# set ge-0/1/3 unit 0 family mpls
user@PCC# set lo0 unit 0 family inet address 192.168.12.1/32
```

2. Enable RSVP on all interfaces of the PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set rsvp interface all
user@PCC# set rsvp interface fxp0.0 disable
```

3. Enable external control of LSPs by the PCEs.

```
[edit protocols]
user@PCC# set mpls lsp-external-controller pccd
```

4. Enable MPLS on all interfaces of the PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set mpls interface all
user@PCC# set mpls interface fxp0.0 disable
```


5. Configure OSPF on all interfaces of the PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set ospf traffic-engineering
user@PCC# set ospf area 0.0.0.0 interface all
user@PCC# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PCC# set ospf interface lo0.0
```

6. Define the PCE group and enable support of PCE-initiated LSPs for the PCE group.

```
[edit protocols]
user@PCC# set protocols pcep pce-group PCEGROUP pce-type active
user@PCC# set protocols pcep pce-group PCEGROUP pce-type stateful
user@PCC# set protocols pcep pce-group PCEGROUP lsp-provisioning
user@PCC# set protocols pcep pce-group PCEGROUP lsp-cleanup-timer 30
```

7. Define the PCEs that connect to the PCC.

```
[edit protocols]
user@PCC# set pcep pce PCE1 destination-ipv4-address 192.168.69.58
user@PCC# set pcep pce PCE1 destination-port 4189
user@PCC# set pcep pce PCE1 pce-group PCEGROUP
user@PCC# set pcep pce PCE2 destination-ipv4-address 192.168.70.65
user@PCC# set pcep pce PCE2 destination-port 4189
user@PCC# set pcep pce PCE2 pce-group PCEGROUP
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PCC# show interfaces
ge-0/1/1 {
  unit 0 {
    family inet {
      address 10.0.102.9/24;
    }
    family iso;
    family mpls;
  }
}
```



```

}
ge-0/1/3 {
  unit 0 {
    family inet {
      address 10.0.112.14/24;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.12.1/32;
    }
  }
}
}

```

user@PCC# **show protocols**

```

rsvp {
  interface all;
}
interface fxp0.0 {
  disable;
}
}
mpls {
  lsp-external-controller pccd;
  interface all;
  interface fxp0.0 {
    disable;
  }
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}
pce-group PCEGROUP {
  pce-type active stateful;
}

```



```
lsp-provisioning;
lsp-cleanup-timer 30;
}
pce PCE1 {
  destination-ipv4-address 192.168.69.58;
  destination-port 4189;
  pce-group PCEGROUP;
}
pce PCE2 {
  destination-ipv4-address 192.168.70.65;
  destination-port 4189;
  pce-group PCEGROUP;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying PCC Status | 1742](#)
- [Verifying PCE1 Status | 1743](#)
- [Verifying the PCE-Initiated LSP Status When the LSP Is Externally Provisioned | 1745](#)

Confirm that the configuration is working properly.

Verifying PCC Status

Purpose

Verify the PCEP session status and LSP summary between the PCC and the connected PCEs.

Action

From operational mode, run the **show path-computation-client status** command.

```
user@PCC# show path-computation-client status
```

Session	Type	Provisioning	Status
PCE1	Stateful Active	On	Up
PCE2	Stateful Active	On	Up


```

LSP Summary
  Total number of LSPs      : 1
  Static LSPs               : 0
  Externally controlled LSPs : 0
  Externally provisioned LSPs : 1/16000 (current/limit)
  Orphaned LSPs            : 0

PCE1 (main)
  Delegated                 : 1
  Externally provisioned    : 1
PCE2
  Delegated                 : 0
  Externally provisioned    : 0

```

Meaning

The output displays the status of the PCEP session between the active stateful PCEs and the PCC. It also displays information about the different types of LSPs on the PCC, and the number of LSPs provisioned by the connected PCEs and delegated to them.

PCE1 is the main active PCE and has one PCE-initiated LSP that has been automatically delegated to it by the PCC.

Verifying PCE1 Status

Purpose

Verify the status of the main active stateful PCE.

Action

From operational mode, run the **show path-computation-client active-pce detail** command.

```
user@PCC# show path-computation-client active-pce
```

```

PCE PCE1
-----
General
  IP address      : 192.168.69.58
  Priority        : 0
  PCE status      : PCE_STATE_UP
  Session type    : PCE_TYPE_STATEFULACTIVE
  LSP provisioning allowed : On
  LSP cleanup timer : 30 [s]
  PCE-mastership  : main

```



```

Max unknown messages      : 5
Keepalives received       : 0
Keepalives sent           : 0
Dead timer                 : 0 [s]
Elapsed as main current    : 1 [s]
Elapsed as main total      : 446380 [s]
Unknown msgs/min rate     : 0
Session failures           : 2198
Corrupted messages        : 0
Delegation timeout set    : 30
Delegation timeout in     : 0 [s]
Delegation failures       : 0
Connection down            : 167092 [s]

```

Counters

PCReqs	Total: 0	last 5min: 0	last hour: 0
PCReps	Total: 0	last 5min: 0	last hour: 0
PCRpts	Total: 5	last 5min: 5	last hour: 5
PCUpdates	Total: 0	last 5min: 0	last hour: 0
PCCreates	Total: 1	last 5min: 1	last hour: 1

Timers

```

Local  Keepalive timer:  30 [s]  Dead timer:  120 [s]  LSP cleanup timer:
30 [s]
Remote Keepalive timer:   0 [s]  Dead timer:    0 [s]  LSP cleanup timer:
- [s]

```

Errors

```

PCErr-recv
PCErr-sent
PCE-PCC-NTFS
PCC-PCE-NTFS

```

Meaning

The output displays information about the current active stateful PCE to which the PCC is connected. The **PCE status** output field indicates the current status of the PCEP session between a PCE and PCC.

For PCE1, the status of the PCEP session is **PCE_STATE_UP**, which indicates that the PCEP session has been established with the PCC.

Verifying the PCE-Initiated LSP Status When the LSP Is Externally Provisioned

Purpose

Verify the status of the PCE-initiated LSP.

Action

From operational mode, run the **show mpls lsp externally-provisioned detail** command.

```
user@PCC# show mpls lsp externally-provisioned detail
```

```
Ingress LSP: 1 sessions

10.0.101.10
  From: 10.0.102.9, State: Up, ActiveRoute: 0, LSPname: lsp15
  ActivePath: path1 (primary)
  Link protection desired
  LSPTYPE: Externally Provisioned, Penultimate hop popping
  LSP Control Status: Externally controlled
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary path1 State: Up
    Priorities: 7 0
    Bandwidth: 8Mbps
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
    10.0.102.10 S 10.0.101.9 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
    10=SoftPreempt 20=Node-ID):
    10.0.102.10 S 10.0.101.9 S
```

Meaning

In the output, the **LSPTYPE** output field shows that the LSP is externally provisioned.

The PCEP session between PCC and PCE1 is up, and the PCC receives the following PCE-initiated LSP parameters:

- ERO (path)—10.0.102.10 and 10.0.101.9
- Bandwidth—8 Mbps
- Priority—7 0 (setup and hold values)

RELATED DOCUMENTATION

[Understanding Static Segment Routing LSP in MPLS Networks | 632](#)

[Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697](#)

[Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE | 1717](#)

[Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1746](#)

[Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs | 1771](#)

Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs

You can configure a Path Computation Client (PCC) with the capability of supporting dynamically created label switched paths (LSPs) from a centralized external path computing entity. A stateful Path Computaiton Element (PCE) can be used to perform external path computation and generate dynamic LSPs when there is an increase in demand.

A PCC creates the PCE-initiated point-to-point LSP using the PCE-provided LSP parameters, or parameters from a pre-configured LSP template when the PCE does not provision the LSP, and automatically delegates the PCE-initiated point-to-point LSP to the respective PCE. As a result, for PCE-initiated LSPs, there is no need for a locally configured LSP on the PCC.

A CLI-controlled LSP, PCE-controlled LSP, and PCE-initiated LSP can coexist with each other on a PCC.

Before you begin:

- Configure the device interfaces.
- Configure MPLS and RSVP-TE.
- Configure OSPF or any other IGP protocol.

To configure PCC to support PCE-initiated point-to-point LSPs, complete the following tasks:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@PCC# edit protocols pcep
```

2. Specify the number of messages per minute that the PCC can receive at maximum.

```
[edit protocols pcep]
user@PCC# set message-rate-limit messages-per-minute
```


- Specify the number of externally provisioned label switched paths (LSPs) over all connected PCEs that the PCC can accept at maximum.

```
[edit protocols pcep]
user@PCC# set max-provisioned-lsps max-count
```

- Specify the unique user defined ID for the connected PCE to configure the PCE parameters.

```
[edit protocols pcep]
user@PCC# edit pce pce-id
```

- Specify the amount of time (in seconds) that the PCC must wait before returning control of LSPs to the routing protocol process after a PCEP session is disconnected.

```
[edit protocols pcep pce pce-id]
user@PCC# set delegation-cleanup-timeout seconds
```

- Specify the IPv4 address of the PCE to connect with.

```
[edit protocols pcep pce pce-id]
user@PCC# set destination-ipv4-address ipv4-address
```

- Specify the TCP port number that the PCE is using

```
[edit protocols pcep pce pce-id]
user@PCC# set destination-port port-number
```

The value can range from 1 through 65535 and the default value is 4189.

- Specify the amount of time (in seconds) that the PCC must wait before deleting any non-delegated PCE-initiated LSPs from the failed PCE after a PCEP session terminates.

```
[edit protocols pcep pce pce-id]
user@PCC# set lsp-cleanup-timer seconds
```

- Configure the PCC to accept SPs that are externally provisioned by connected PCEs. By default, the PCC rejects PCE-initiated LSPs.


```
[edit protocols pcep pce pce-id]
user@PCC# set lsp-provisioning
```

10. Specify the number of unknown messages per minute that the PCC can receive at maximum after which the PCEP session is closed.

```
[edit protocols pcep pce pce-id]
user@PCC# set max-unknown-messages messages-per-minute
```

The value can range from 1 through 16384, and the default value is 0 (disabled or no limit).

11. Specify the number of unknown requests per minute that the PCC can receive at maximum after which the PCEP session is terminated.

```
[edit protocols pcep pce pce-id]
user@PCC# set max-unknown-requests requests-per-minute
```

The value can range from 0 through 16384, and the default value is 5. A value of 0 disables this statement.

12. Configure the PCE type.

```
[edit protocols pcep pce pce-id]
user@PCC# set pce-type active stateful
```

13. Specify the amount of time (in seconds) that the PCC must wait for a reply before resending a request.

```
[edit protocols pcep pce pce-id]
user@PCC# set request-timer seconds
```

The value can range from 0 through 65535 seconds.

14. Verify and commit the configuration.

```
user@PCC# show
user@PCC# commit
```


Sample Output

```
[edit]
user@PCC# edit protocols pcep

[edit protocols pcep]
user@PCC# set message-rate-limit 50

[edit protocols pcep]
user@PCC# set max-provisioned-lsps 16000

[edit protocols pcep]
user@PCC# edit pce PCE

[edit protocols pcep pce PCE]
user@PCC# set delegation-cleanup-timeout 20

[edit protocols pcep pce PCE]
user@PCC# set destination-ipv4-address 192.168.69.58

[edit protocols pcep pce PCE]
user@PCC# set destination-port 4189

[edit protocols pcep pce PCE]
user@PCC# set lsp-cleanup-timer 50

[edit protocols pcep pce PCE]
user@PCC# set lsp-provisioning

[edit protocols pcep pce PCE]
user@PCC# set max-unknown-messages 5

[edit protocols pcep pce PCE]
user@PCC# set max-unknown-requests 5

[edit protocols pcep pce PCE]
user@PCC# set request-timer 50

[edit protocols pcep pce PCE]
user@PCC# up

[edit protocols pcep]
user@PCC# show
message-rate-limit 50;
max-provisioned-lsps 16000;
```



```
pce PCE {
    destination-ipv4-address 192.168.69.58;
    destination-port 4189;
    lsp-provisioning;
    lsp-cleanup-timer 50;
    request-timer 50;
    max-unknown-requests 5;
    max-unknown-messages 5;
    delegation-cleanup-timeout 20;
}

[edit protocols pcep]
user@PCC# commit
commit complete
```

RELATED DOCUMENTATION

[Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697](#)

[Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1734](#)

[pcep | 3072](#)

[Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs | 1771](#)

Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Controlled Point-to-Multipoint LSPs

IN THIS SECTION

- [Requirements | 1751](#)
- [Overview | 1751](#)
- [Configuration | 1752](#)
- [Verification | 1766](#)

This example shows how to configure the Path Computation Client (PCC) with the capability of reporting point-to-multipoint traffic engineered label-switched paths (TE LSPs) to a Path Computation Element (PCE).

Requirements

This example uses the following hardware and software components:

- Three routers that can be a combination of ACX Series, M Series, MX Series, or T Series routers.
- One virtual machine configured with Virtual Route Reflector (VRR) feature.
- A TCP connection to an external stateful PCE from the VRR.
- Junos OS Release 16.1 or later running on the PCC.

Before you begin:

- Configure the device interfaces.
- Configure MPLS and RSVP-TE.
- Configure OSPF or any other IGP protocol.

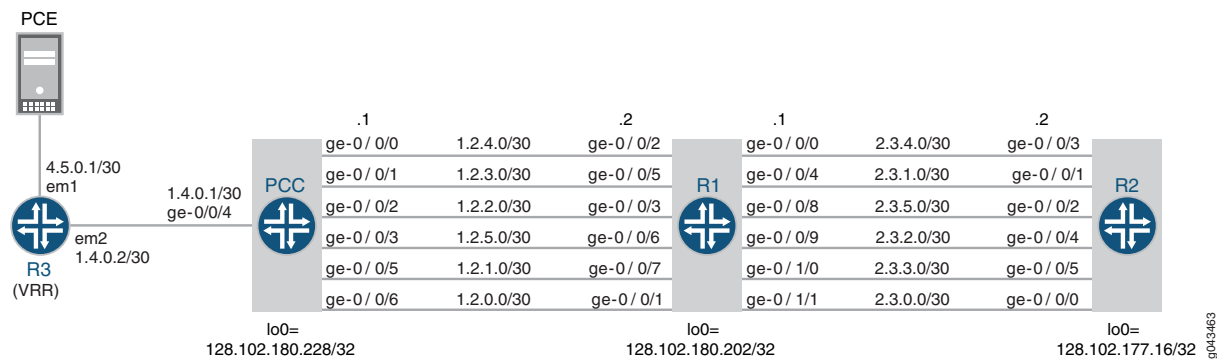
Overview

After a PCEP session is established between a PCE and a PCC, the PCC reports all the LSPs in the system to the PCE for LSP state synchronization. This includes PCC-controlled, PCE-delegated, and PCE-initiated point-to-point LSPs. Starting with Junos OS Release 15.1F6 and 16.1R1, this capability is extended to report point-to-multipoint LSPs as well.

By default, PCE control of point-to-multipoint LSPs is not supported on a PCC. To add this capability, include the **p2mp-lsp-report-capability** statement at the **[edit protocols pcep pce *pce-name*]** or **[edit protocols pcep pce-group *group-id*]** hierarchy levels.

Topology

Figure 142: Example PCE-Controlled Point-to-Multipoint LSPs



In this example, PCC is the ingress router, Router R1 is the transit router, and Router R2 is the egress router. PCC is connected to a Virtual Route Reflector (VRR) that is connected to a PCE. There are many point-to-multipoint interfaces between PCC, Router R1, and Router R2.

The reporting of point-to-multipoint LSPs is executed as follows:

1. If Router PCC is configured with point-to-point and point-to-multipoint LSPs without the support for point-to-multipoint reporting capability, only the point-to-point LSPs are reported to the connected PCE. By default, a PCC does not support point-to-multipoint LSP reporting capability.
2. When Router PCC is configured with point-to-multipoint LSP reporting capability, PCC first advertises this capability to PCE through a report message.
3. By default, a PCE provides support for point-to-multipoint LSP capability. On receiving the PCC's advertisement for point-to-multipoint LSP capability, the PCE in return advertises its capability to the PCC.
4. On receiving the PCE's advertisement of the point-to-multipoint capability, PCC reports all branches of point-to-multipoint LSPs to the PCE using the update message.
5. Once all the LSPs are reported to the PCE, LSP state is synchronized between the PCE and PCC.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

PCC

```

set interfaces ge-0/0/0 unit 0 family inet address 1.2.4.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 1.2.3.1/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 1.2.2.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 1.2.5.1/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 unit 0 family inet address 1.4.0.1/30
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces ge-0/0/5 unit 0 family inet address 1.2.1.1/30
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family inet address 1.2.0.1/30
set interfaces ge-0/0/6 unit 0 family mpls
set routing-options autonomous-system 100
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls lsp-external-controller pccd pce-controlled-lsp pcc_delegated_no_cspf_*
    label-switched-path-template lsp_template_no_cspf
set protocols mpls lsp-external-controller pccd pce-controlled-lsp pce_initiated_no_ero_no_cspf_*
    label-switched-path-template lsp_template_no_cspf
set protocols mpls lsp-external-controller pccd pce-controlled-lsp pce_initiated_loose_ero_no_cspf_*
    label-switched-path-template lsp_template_no_cspf
set protocols mpls traffic-engineering database import policy TE
set protocols mpls admin-groups violet 1
set protocols mpls admin-groups indigo 2
set protocols mpls admin-groups blue 3
set protocols mpls admin-groups green 4
set protocols mpls admin-groups yellow 5
set protocols mpls admin-groups orange 6
set protocols mpls label-switched-path lsp_template_no_cspf template
set protocols mpls label-switched-path lsp_template_no_cspf no-cspf
set protocols mpls label-switched-path lsp1-pcc to 128.102.177.16
set protocols mpls label-switched-path lsp2-pcc to 128.102.177.16
set protocols mpls label-switched-path lsp2-pcc lsp-external-controller pccd
set protocols mpls path loose-path 1.2.3.2 loose
set protocols mpls path strict-path 1.2.3.2 strict
set protocols mpls path strict-path 2.3.3.2 strict
set protocols mpls path path-B
set protocols mpls path path-C
set protocols mpls interface all

```



```

set protocols mpls interface ge-0/0/6.0 admin-group violet
set protocols mpls interface ge-0/0/5.0 admin-group indigo
set protocols mpls interface ge-0/0/2.0 admin-group blue
set protocols mpls interface ge-0/0/1.0 admin-group green
set protocols mpls interface ge-0/0/0.0 admin-group yellow
set protocols mpls interface ge-0/0/3.0 admin-group orange
set protocols mpls interface fxp0.0 disable
set protocols bgp group northstar type internal
set protocols bgp group northstar local-address 128.102.180.228
set protocols bgp group northstar family traffic-engineering unicast
set protocols bgp group northstar export TE
set protocols bgp group northstar neighbor 128.102.180.215
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p
set protocols pcep pce pce1 local-address 10.102.180.228
set protocols pcep pce pce1 destination-ipv4-address 10.102.180.246
set protocols pcep pce pce1 destination-port 4189
set protocols pcep pce pce1 pce-type active
set protocols pcep pce pce1 pce-type stateful
set protocols pcep pce pce1 lsp-provisioning
set protocols pcep pce pce1 lsp-cleanup-timer 0
set protocols pcep pce pce1 delegation-cleanup-timeout 60
set protocols pcep pce pce1 p2mp-lsp-report-capability
set policy-options policy-statement TE term 1 from family traffic-engineering
set policy-options policy-statement TE term 1 then accept

```

R1

```

set interfaces ge-0/0/0 unit 0 family inet address 2.3.4.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 1.2.0.2/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 1.2.4.2/30

```



```
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 1.2.2.2/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 unit 0 family inet address 2.3.1.1/30
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces ge-0/0/5 unit 0 family inet address 1.2.3.2/30
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family inet address 1.2.5.2/30
set interfaces ge-0/0/6 unit 0 family mpls
set interfaces ge-0/0/7 unit 0 family inet address 1.2.1.2/30
set interfaces ge-0/0/7 unit 0 family mpls
set interfaces ge-0/0/8 unit 0 family inet address 2.3.5.1/30
set interfaces ge-0/0/8 unit 0 family mpls
set interfaces ge-0/0/9 unit 0 family inet address 2.3.2.1/30
set interfaces ge-0/0/9 unit 0 family mpls
set interfaces ge-0/1/0 unit 0 family inet address 2.3.3.1/30
set interfaces ge-0/1/0 unit 0 family mpls
set interfaces ge-0/1/1 unit 0 family inet address 2.3.0.1/30
set interfaces ge-0/1/1 unit 0 family mpls
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls admin-groups violet 1
set protocols mpls admin-groups indigo 2
set protocols mpls admin-groups blue 3
set protocols mpls admin-groups green 4
set protocols mpls admin-groups yellow 5
set protocols mpls admin-groups orange 6
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/1.0 admin-group violet
set protocols mpls interface ge-0/0/7.0 admin-group indigo
set protocols mpls interface ge-0/0/3.0 admin-group blue
set protocols mpls interface ge-0/0/5.0 admin-group green
set protocols mpls interface ge-0/0/2.0 admin-group yellow
set protocols mpls interface ge-0/0/6.0 admin-group orange
set protocols mpls interface ge-0/1/1.0 admin-group violet
set protocols mpls interface ge-0/0/4.0 admin-group indigo
set protocols mpls interface ge-0/0/9.0 admin-group blue
set protocols mpls interface ge-0/1/0.0 admin-group green
set protocols mpls interface ge-0/0/0.0 admin-group yellow
set protocols mpls interface ge-0/0/8.0 admin-group orange
set protocols ospf traffic-engineering
```



```

set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/7.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols ospf area 0.0.0.0 interface ge-0/1/1.0

```

R2

```

set interfaces ge-0/0/0 unit 0 family inet address 2.3.0.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 2.3.1.2/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 2.3.5.2/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 2.3.4.2/30
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 unit 0 family inet address 2.3.2.2/30
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces ge-0/0/5 unit 0 family inet address 2.3.3.2/30
set interfaces ge-0/0/5 unit 0 family mpls
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls admin-groups violet 1
set protocols mpls admin-groups indigo 2
set protocols mpls admin-groups blue 3
set protocols mpls admin-groups green 4
set protocols mpls admin-groups yellow 5
set protocols mpls admin-groups orange 6
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls interface ge-0/0/0.0 admin-group violet
set protocols mpls interface ge-0/0/1.0 admin-group indigo
set protocols mpls interface ge-0/0/4.0 admin-group blue
set protocols mpls interface ge-0/0/5.0 admin-group green
set protocols mpls interface ge-0/0/3.0 admin-group yellow
set protocols mpls interface ge-0/0/2.0 admin-group orange
set protocols ospf traffic-engineering

```



```

set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

R3

```

set interfaces em0 unit 0 family inet address 10.102.180.215/19
set interfaces em1 unit 0 family inet address 4.5.0.1/30
set interfaces em2 unit 0 family inet address 1.4.0.2/30
set interfaces em2 unit 0 family mpls
set routing-options router-id 128.102.180.215
set routing-options autonomous-system 100
set protocols topology-export
set protocols rsvp interface all
set protocols mpls lsp-external-controller pccd
set protocols mpls traffic-engineering database import igp-topology
set protocols mpls traffic-engineering database import policy TE
set protocols mpls interface all
set protocols bgp group northstar type internal
set protocols bgp group northstar local-address 128.102.180.215
set protocols bgp group northstar family traffic-engineering unicast
set protocols bgp group northstar neighbor 128.102.180.228
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface em2.0 interface-type p2p
set policy-options policy-statement TE from family traffic-engineering
set policy-options policy-statement TE then accept

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the PCC router:

1. Configure the interfaces of Router PCC. To enable MPLS, include the protocol family on the interface so that the interface does not discard incoming MPLS traffic.


```
[edit interfaces]
user@PCC# set ge-0/0/0 unit 0 family inet address 1.2.4.1/30
user@PCC# set ge-0/0/0 unit 0 family mpls
user@PCC# set ge-0/0/1 unit 0 family inet address 1.2.3.1/30
user@PCC# set ge-0/0/1 unit 0 family mpls
user@PCC# set ge-0/0/2 unit 0 family inet address 1.2.2.1/30
user@PCC# set ge-0/0/2 unit 0 family mpls
user@PCC# set ge-0/0/3 unit 0 family inet address 1.2.5.1/30
user@PCC# set ge-0/0/3 unit 0 family mpls
user@PCC# set ge-0/0/4 unit 0 family inet address 1.4.0.1/30
user@PCC# set ge-0/0/4 unit 0 family mpls
user@PCC# set ge-0/0/5 unit 0 family inet address 1.2.1.1/30
user@PCC# set ge-0/0/5 unit 0 family mpls
user@PCC# set ge-0/0/6 unit 0 family inet address 1.2.0.1/30
user@PCC# set ge-0/0/6 unit 0 family mpls
```

2. Configure the autonomous system number for Router PCC.

```
[edit routing-options]
user@PCC# set autonomous-system 100
```

3. Enable RSVP on all interfaces of Router PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set rsvp interface all
user@PCC# set rsvp interface fxp0.0 disable
```

4. Enable MPLS on all the interfaces of Router PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set mpls interface all
user@PCC# set mpls interface fxp0.0 disable
```

5. Configure a dynamic LSP and disable automatic path computation for the LSP.

```
[edit protocols]
user@PCC# set mpls label-switched-path lsp_template_no_cspf template
user@PCC# set mpls label-switched-path lsp_template_no_cspf no-cspf
```

6. Configure point-to-multipoint LSPs and define external path computing entity for the LSP.


```
[edit protocols]
user@PCC# set mpls label-switched-path lsp1-pcc to 128.102.177.16
user@PCC# set mpls label-switched-path lsp2-pcc to 128.102.177.16
user@PCC# set mpls label-switched-path lsp2-pcc lsp-external-controller pccd
```

7. Enable external path computing for the MPLS LSPs and assign a template for externally provisioned LSPs.

```
[edit protocols]
user@PCC# set mpls lsp-external-controller pccd pce-controlled-lsp pcc_delegated_no_cspf_*
    label-switched-path-template lsp_template_no_cspf
user@PCC# set mpls lsp-external-controller pccd pce-controlled-lsp pce_initiated_no_ero_no_cspf_*
    label-switched-path-template lsp_template_no_cspf
user@PCC# set mpls lsp-external-controller pccd pce-controlled-lsp pce_initiated_loose_ero_no_cspf_*
    label-switched-path-template lsp_template_no_cspf
```

8. Configure the LSPs that have local control and are overridden by the PCE-provided LSP parameters.

```
[edit protocols]
user@PCC# set mpls path loose-path 1.2.3.2 loose
user@PCC# set mpls path strict-path 1.2.3.2 strict
user@PCC# set mpls path strict-path 2.3.3.2 strict
user@PCC# set mpls path path-B
user@PCC# set mpls path path-C
```

9. Configure MPLS administrative group policies for constrained-path LSP computation.

```
[edit protocols]
user@PCC# set mpls admin-groups violet 1
user@PCC# set mpls admin-groups indigo 2
user@PCC# set mpls admin-groups blue 3
user@PCC# set mpls admin-groups green 4
user@PCC# set mpls admin-groups yellow 5
user@PCC# set mpls admin-groups orange 6
```

10. Assign the configured administrative group policies to Router PCC interfaces.

```
[edit protocols]
user@PCC# set mpls interface ge-0/0/6.0 admin-group violet
user@PCC# set mpls interface ge-0/0/5.0 admin-group indigo
```



```

user@PCC# set mpls interface ge-0/0/2.0 admin-group blue
user@PCC# set mpls interface ge-0/0/1.0 admin-group green
user@PCC# set mpls interface ge-0/0/0.0 admin-group yellow
user@PCC# set mpls interface ge-0/0/3.0 admin-group orange

```

11. Configure a traffic engineering database (TED) import policy.

```

[edit protocols]
user@PCC# set mpls traffic-engineering database import policy TE

```

12. Configure a BGP internal group.

```

[edit protocols]
user@PCC# set bgp group northstar type internal
user@PCC# set bgp group northstar local-address 128.102.180.228
user@PCC# set bgp group northstar neighbor 128.102.180.215

```

13. Configure traffic engineering for BGP and assign the export policy.

```

[edit protocols]
user@PCC# set bgp group northstar family traffic-engineering unicast
user@PCC# set bgp group northstar export TE

```

14. Configure OSPF area 0 on all the point-to-multipoint interfaces of Router PCC.

```

[edit protocols]
user@PCC# set ospf area 0.0.0.0 interface lo0.0
user@PCC# set ospf area 0.0.0.0 interface ge-0/0/6.0
user@PCC# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@PCC# set ospf area 0.0.0.0 interface ge-0/0/2.0
user@PCC# set ospf area 0.0.0.0 interface ge-0/0/1.0
user@PCC# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@PCC# set ospf area 0.0.0.0 interface ge-0/0/3.0

```

15. Configure OSPF area 0 on the point-to-point interface of Router PCC.

```

[edit protocols]
user@PCC# set ospf area 0.0.0.0 interface ge-0/0/4.0 interface-type p2p

```


16. Enable traffic engineering for OSPF.

```
[edit protocols]
user@PCC# set ospf traffic-engineering
```

17. Define the PCE that Router PCC connects to, and configure the the PCE parameters.

```
[edit protocols]
user@PCC# set pcep pce pce1 local-address 10.102.180.228
user@PCC# set pcep pce pce1 destination-ipv4-address 10.102.180.246
user@PCC# set pcep pce pce1 destination-port 4189
user@PCC# set pcep pce pce1 pce-type active
user@PCC# set pcep pce pce1 pce-type stateful
user@PCC# set pcep pce pce1 lsp-provisioning
user@PCC# set pcep pce pce1 lsp-cleanup-timer 0
user@PCC# set pcep pce pce1 delegation-cleanup-timeout 60
```

18. Configure Router PCC to enable point-to-multipoint LSP capability for external path computing.

```
[edit protocols]
set pcep pce pce1 p2mp-lsp-report-capability
```

19. Configure the traffic engineering policy.

```
[edit policy-options]
user@PCC# set policy-statement TE term 1 from family traffic-engineering
user@PCC# set policy-statement TE term 1 then accept
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PCC# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 1.2.4.1/30;
    }
  }
}
```



```
        family mpls;
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 1.2.3.1/30;
        }
        family mpls;
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 1.2.2.1/30;
        }
        family mpls;
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 1.2.5.1/30;
        }
        family mpls;
    }
}
ge-0/0/4 {
    unit 0 {
        family inet {
            address 1.4.0.1/30;
        }
        family mpls;
    }
}
ge-0/0/5 {
    unit 0 {
        family inet {
            address 1.2.1.1/30;
        }
        family mpls;
    }
}
ge-0/0/6 {
```



```

unit 0 {
    family inet {
        address 1.2.0.1/30;
    }
    family mpls;
}
}

```

user@PCC# **show protocols**

```

rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    lsp-external-controller pccd {
        pce-controlled-lsp pcc_delegated_no_cspf_* {
            label-switched-path-template {
                lsp_template_no_cspf;
            }
        }
        pce-controlled-lsp pce_initiated_no_ero_no_cspf_* {
            label-switched-path-template {
                lsp_template_no_cspf;
            }
        }
        pce-controlled-lsp pce_initiated_loose_ero_no_cspf_* {
            label-switched-path-template {
                lsp_template_no_cspf;
            }
        }
    }
}
traffic-engineering {
    database {
        import {
            policy TE;
        }
    }
}
admin-groups {
    violet 1;
    indigo 2;
    blue 3;
}

```



```

    green 4;
    yellow 5;
    orange 6;
}
label-switched-path lsp_template_no_cspf {
    template;
    no-cspf;
}
label-switched-path lsp1-pcc {
    to 128.102.177.16;
}
label-switched-path lsp2-pcc {
    to 128.102.177.16;
    lsp-external-controller pccd;
}
path loose-path {
    1.2.3.2 loose;
}
path strict-path {
    1.2.3.2 strict;
    2.3.3.2 strict;
}
path path-B;
path path-C;
interface all;
interface ge-0/0/6.0 {
    admin-group violet;
}
interface ge-0/0/5.0 {
    admin-group indigo;
}
interface ge-0/0/2.0 {
    admin-group blue;
}
interface ge-0/0/1.0 {
    admin-group green;
}
interface ge-0/0/0.0 {
    admin-group yellow;
}
interface ge-0/0/3.0 {
    admin-group orange;
}
interface fxp0.0 {

```



```

        disable;
    }
}
bgp {
    group northstar {
        type internal;
        local-address 128.102.180.228;
        family traffic-engineering {
            unicast;
        }
        export TE;
        neighbor 128.102.180.215;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/6.0;
        interface ge-0/0/5.0;
        interface ge-0/0/2.0;
        interface ge-0/0/1.0;
        interface ge-0/0/0.0;
        interface ge-0/0/3.0;
        interface ge-0/0/4.0 {
            interface-type p2p;
        }
    }
}
pcep {
    pce pce1 {
        local-address 10.102.180.228;
        destination-ipv4-address 10.102.180.246;
        destination-port 4189;
        pce-type active stateful;
        lsp-provisioning;
        lsp-cleanup-timer 0;
        delegation-cleanup-timeout 60;
        p2mp-lsp-report-capability;
    }
}
}

```


Verification

IN THIS SECTION

- [Verifying LSP Configuration on the PCC | 1766](#)
- [Verifying PCE Configuration on the PCC | 1770](#)

Confirm that the configuration is working properly.

Verifying LSP Configuration on the PCC

Purpose

Verify the LSP type and running state of the point-to-multipoint LSP.

Action

From operational mode, run the **show mpls lsp extensive** command.

```
user@PCC> show mpls lsp extensive
```

```
Ingress LSP: 2 sessions

128.102.177.16
  From: 128.102.180.228, State: Up, ActiveRoute: 0, LSPname: lsp1-pcc
  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: 2)
  1.2.1.2 S 2.3.0.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
      1.2.1.2 2.3.0.2
    6 Jul 12 14:44:10.620 Selected as active path
    5 Jul 12 14:44:10.617 Record Route: 1.2.1.2 2.3.0.2
    4 Jul 12 14:44:10.615 Up
    3 Jul 12 14:44:10.175 Originate Call
    2 Jul 12 14:44:10.174 CSPF: computation result accepted 1.2.1.2 2.3.0.2
    1 Jul 12 14:43:41.442 CSPF failed: no route toward 128.102.177.16[2 times]
```


Created: Tue Jul 12 14:42:43 2016

128.102.177.16

From: 128.102.180.228, State: Up, ActiveRoute: 0, **LSPname: lsp2-pcc**

ActivePath: (primary)

LSPtype: Externally controlled - static configured, Penultimate hop popping

LSP Control Status: Externally controlled

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

*Primary State: Up

Priorities: 7 0

External Path CSPF Status: external

SmartOptimizeTimer: 180

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):

1.2.4.2 2.3.0.2

50 Jul 12 14:50:14.699 EXTCTRL LSP: Sent Path computation request and LSP status

49 Jul 12 14:50:14.698 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0

48 Jul 12 14:49:27.859 EXTCTRL LSP: Sent Path computation request and LSP status

47 Jul 12 14:49:27.859 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0

46 Jul 12 14:49:27.858 EXTCTRL LSP: Sent Path computation request and LSP status

45 Jul 12 14:49:27.858 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0

44 Jul 12 14:49:27.858 EXTCTRL_LSP: Control status became external

43 Jul 12 14:49:03.746 EXTCTRL_LSP: Control status became local

42 Jul 12 14:46:52.367 EXTCTRL LSP: Sent Path computation request and LSP status

41 Jul 12 14:46:52.367 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0

40 Jul 12 14:46:52.367 EXTCTRL LSP: Sent Path computation request and LSP status

39 Jul 12 14:46:52.366 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0

38 Jul 12 14:46:52.366 EXTCTRL_LSP: Control status became external


```
37 Jul 12 14:46:41.584 Selected as active path
36 Jul 12 14:46:41.565 Record Route: 1.2.4.2 2.3.0.2
35 Jul 12 14:46:41.565 Up
34 Jul 12 14:46:41.374 EXTCTRL_LSP: Applying local parameters with this
signalling attempt
33 Jul 12 14:46:41.374 Originate Call
32 Jul 12 14:46:41.374 CSPF: computation result accepted 1.2.4.2 2.3.0.2
31 Jul 12 14:46:28.254 EXTCTRL_LSP: Control status became local
30 Jul 12 14:46:12.494 EXTCTRL LSP: Sent Path computation request and LSP status

29 Jul 12 14:46:12.494 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
28 Jul 12 14:45:43.164 EXTCTRL LSP: Sent Path computation request and LSP status

27 Jul 12 14:45:43.164 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
26 Jul 12 14:45:13.424 EXTCTRL LSP: Sent Path computation request and LSP status

25 Jul 12 14:45:13.423 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
24 Jul 12 14:44:44.774 EXTCTRL LSP: Sent Path computation request and LSP status

23 Jul 12 14:44:44.773 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
22 Jul 12 14:44:15.053 EXTCTRL LSP: Sent Path computation request and LSP status

21 Jul 12 14:44:15.053 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
20 Jul 12 14:43:45.705 EXTCTRL LSP: Sent Path computation request and LSP status

19 Jul 12 14:43:45.705 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
18 Jul 12 14:43:45.705 EXTCTRL LSP: Sent Path computation request and LSP status

17 Jul 12 14:43:45.705 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
16 Jul 12 14:43:45.705 EXTCTRL_LSP: Control status became external
```



```

15 Jul 12 14:43:42.398 CSPF failed: no route toward 128.102.177.16
14 Jul 12 14:43:13.009 EXTCTRL_LSP: Control status became local
13 Jul 12 14:43:13.009 EXTCTRL LSP: Sent Path computation request and LSP status

12 Jul 12 14:43:13.008 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
11 Jul 12 14:42:43.343 EXTCTRL LSP: Sent Path computation request and LSP status

10 Jul 12 14:42:43.343 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
9 Jul 12 14:42:43.343 EXTCTRL LSP: Sent Path computation request and LSP status

8 Jul 12 14:42:43.343 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
7 Jul 12 14:42:43.342 EXTCTRL LSP: Sent Path computation request and LSP status

6 Jul 12 14:42:43.342 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
5 Jul 12 14:42:43.341 EXTCTRL_LSP: Control status became external
4 Jul 12 14:42:43.337 EXTCTRL_LSP: Control status became local
3 Jul 12 14:42:43.323 EXTCTRL LSP: Sent Path computation request and LSP status

2 Jul 12 14:42:43.323 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 0 req BW 0 admin group(exclude 0 include any 0 include all 0) priority
setup 7 hold 0
1 Jul 12 14:42:43.258 EXTCTRL LSP: Awaiting external controller connection
Created: Tue Jul 12 14:42:43 2016
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The output displays the lsp2-pcc LSP as the PCE-controlled LSP.

Verifying PCE Configuration on the PCC

Purpose

Verify the PCE parameters configuration and PCE state.

Action

From operational mode, run the **show path-computation-client active-pce** command.

```
user@PCC> show path-computation-client active-pce
```

```
PCE pce1
-----
General
  PCE IP address      : 10.102.180.246
  Local IP address    : 10.102.180.228
  Priority             : 0
  PCE status          : PCE_STATE_UP
  Session type        : PCE_TYPE_STATEFULACTIVE
  LSP provisioning allowed : On
  P2MP LSP report allowed : On
  P2MP LSP update allowed : Off
  P2MP LSP init allowed  : Off
  PCE-mastership       : main

Counters
  PCReqs              Total: 0          last 5min: 0          last hour: 0
  PCReps              Total: 0          last 5min: 0          last hour: 0
  PCRpts              Total: 12         last 5min: 0          last hour: 12
  PCUpdates           Total: 1          last 5min: 0          last hour: 1
  PCCreates           Total: 0          last 5min: 0          last hour: 0

Timers
  Local  Keepalive timer: 30 [s]  Dead timer: 120 [s]  LSP cleanup timer:
0 [s]
  Remote Keepalive timer: 30 [s]  Dead timer: 120 [s]  LSP cleanup timer:
0 [s]

Errors
  PCErr-recv
```



```

PCErr-sent
    Type: 1          Value: 2          Count: 1
PCE-PCC-NTFS
PCC-PCE-NTFS

```

Meaning

The output displays the active PCE that Router PCC is connected to, and the pce1 PCE parameters and state.

RELATED DOCUMENTATION

[Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697](#)

Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs

IN THIS SECTION

- [Benefits of PCE-Initiated Point-to-Multipoint LSPs | 1772](#)
- [Signaling of PCE-Initiated Point-to-Multipoint LSPs | 1772](#)
- [Behavior of PCE-Initiated Point-to-Multipoint LSPs After PCEP Session Failure | 1772](#)
- [Configuring PCE-Initiated Point-to-Multipoint LSP Capability | 1773](#)
- [Supported and Unsupported Features for PCE-Initiated Point-to-Multipoint LSPs | 1773](#)
- [Mapping PCE-initiated Point-To-Multipoint LSPs to MVPN | 1773](#)

With the introduction of point-to-multipoint PCE-initiated LSPs, a PCE can initiate and provision a point-to-multipoint LSP dynamically without the need for local LSP configuration on the PCC. This enables the PCE to control the timing and sequence of the point-to-multipoint path computations within and across Path Computation Element Protocol (PCEP) sessions, thereby creating a dynamic network that is centrally controlled and deployed.

Benefits of PCE-Initiated Point-to-Multipoint LSPs

Meets the requirements of point-to-multipoint traffic engineering LSP placement in response to application demands through dynamic creation and tear down of point-to-multipoint LSPs, thereby creating a dynamic network that is centrally controlled and deployed.

Signaling of PCE-Initiated Point-to-Multipoint LSPs

The signaling of PCE-initiated point-to-multipoint LSPs is as follows:

- **When a new branch is added (Grafting)**—Only the new branch sub-LSP is signaled and does not result in re-signaling of the entire point-to-multipoint tree.

If any topology changes occurred before provisioning of the new sub-LSP, then the Path Computation Server (PCS) re-computes the entire point-to-multipoint tree and updates the point-to-multipoint LSP using a PC update message.

- **When a branch is deleted (Pruning)**—The deleted branch sub-LSP is torn down and does not result in re-signaling of the entire point-to-multipoint tree.
- **When a branch sub-LSP parameter is changed**—Change in sub-LSP parameters, such as Explicit Route Object (ERO), bandwidth, or priority, can happen either because of optimization, or on user request. If there is a re-signaling request for a sub-LSP, the entire point-to-multipoint tree is re-signaled, and then the switchover to the new instance happens once the new instances of all the branches are up.
- **When a branch sub-LSP path fails**—An error is reported to the PCS for the failed branch sub-LSP. On receiving the new ERO from the PCS, the entire point-to-multipoint tree is re-signaled along with the failed branch sub-LSP, and the switchover to the new instance happens in a make-before-break (MBB) fashion.

Behavior of PCE-Initiated Point-to-Multipoint LSPs After PCEP Session Failure

When a PCEP session fails, the PCE-initiated point-to-multipoint LSPs are orphaned until the expiration of the **state timeout** timer. After the **state timeout** timer expires, the PCE-initiated LSPs are cleaned up.

To obtain control of a PCE-initiated point-to-multipoint LSP after a PCEP session failure, the primary or secondary PCE sends a **PCInitiate** message before the **state timeout** timer expires.

Configuring PCE-Initiated Point-to-Multipoint LSP Capability

By default, the creation and provisioning of point-to-multipoint LSPs by a PCE is not supported on a PCC. To enable this capability, include the **p2mp-lsp-init-capability** and **p2mp-lsp-update-capability** statements at the **[edit protocols pcep pce pce-name]** or **[edit protocols pcep pce-group group-id]** hierarchy levels.

The **p2mp-lsp-init-capability** statement provides the capability to provision point-to-multipoint RSVP-TE LSPs by a PCE. The **p2mp-lsp-update-capability** statement provides the capability to update point-to-multipoint RSVP-TE LSP parameters by a PCE.

Supported and Unsupported Features for PCE-Initiated Point-to-Multipoint LSPs

The following features are supported with PCE-initiated point-to-multipoint LSPs:

- Partial compliance with the Internet draft draft-ietf-pce-stateful-pce-p2mp (expires October 2018), *Path Computation Element (PCE) Protocol Extensions for Stateful PCE usage for Point-to-Multipoint Traffic Engineering Label Switched Paths*.

The following features are not supported with PCE-initiated point-to-multipoint LSPs:

- Delegation of point-to-multipoint locally controlled LSP.
- LSP control delegation.
- Interior gateway protocol (IGP) extension for PCE discovery within an IGP routing domain.
- Request/response messaging.
- Direct movement of branch sub-LSP from one point-to-multipoint tree to another.

The same can be achieved by deleting a branch sub-LSP from the first point-to-multipoint tree and re-adding it to another after the **PCReport** message indicates LSP removal from the device.

- IPv6 is not supported.
- SERO based signalling is not supported.
- Empty-ERO feature is not supported.
- Link protection is not supported.

Mapping PCE-initiated Point-To-Multipoint LSPs to MVPN

You can associate a single or range of MVPN multicast flows (S,G) to a dynamically created PCE-initiated point-to-multipoint label-switched path (LSP). You can specify only selective types of flows for this feature to work. This includes:

- Route distinguisher (RD) which is mapped to the MVPN routing-instance.

- (S,G) which is the source of a multicast packet and destination multicast group address. This is used to filter incoming traffic for mapping it to the tunnel.
- Point-to-multipoint LSP that is used to send traffic that matches the above-mentioned flow specification.

For more details, see Internet draft draft-ietf-pce-pcep-flowspec-05 (expires February 16, 2020) *PCEP Extension for Flow Specification*.

The current implementation of this feature does not implement the following sections of the draft:

- Section 3.1.2—Advertising PCE capabilities in IGP
- Section 3.2—PCReq and PCRep message
- Section 7—Most of the flow specifications, except route distinguisher and IPv4 multicast flow specifications, are not supported.

To enable the mapping of PCE-initiated point-to-multipoint LSPs to MVPN:

- Include the **pce_traffic_steering** statement at the [edit protocols pcep pce pce-id] hierarchy level to indicate the support for flow specification capability (also called traffic steering) by the PCC.
- Include the **external-controller** statement at the [edit routing-instances routing-instance-name provider-tunnel] hierarchy level.

The presence of **external-controller** in the provider-tunnel configuration for MVPN indicates that the point-to-multipoint LSP and (S,G) for this MVPN instance can be provided by the external controller. This enables the external controller to dynamically configure (S,G) and point-to-multipoint LSP for MVPN.

Take the following into consideration for mapping of PCE-initiated point-to-multipoint LSPs to MVPN:

- If you do not enable the **external-controller pccd** statement for a particular MVPN instance, then the PCCD process does not configure (S,G) dynamically.
- If you disable the **external-controller pccd** configuration from the CLI, then the dynamically learned multicast flows (S,G) for that particular MVPN instance is deleted and reported to the external controller.
- When (S,G) is already configured from the CLI, the PCC cannot configure (S,G) dynamically as local configuration has a higher priority.
- If any particular (S,G) is learned from the external controller dynamically and then you configure the same (S,G) for the same MVPN instance, then the dynamically learned (S,G) is deleted and reported to the external controller through the PCC.
- If the routing protocol process reboots, then the PCCD process reconfigures all the (S,G) again.
- If the PCCD process reboots, then MVPN reports all PCCD configured (S,G) to the external controller.
- If user enables **external-controller pccd** for a particular MVPN instance, then MVPN requests the PCCD process to configure (S,G), if any present.
- If there are major configuration changes to a particular MVPN instance, then MVPN requests the PCCD process to reconfigure all (S,G) for that particular MVPN instance.

- All flow specifications associated with any PCE-initiated point-to-multipoint LSP must have the same RD. During PC initiation if all flow specifications do not have the same RD, then the PC initiate message is dropped with an error.
- You can associate a point-to-multipoint LSP only with selective type of flow specifications, otherwise the PC initiate message is dropped with an error.
- During PC update if all flow specifications do not have same the RD either due to a new flow specification addition, or due to existing flow specification update, then the PCC drops the update message.
- During PC update if all flow specifications do not meet the selective condition either due to new flow specification addition, or due to existing flow specifications update, then the PCC drops the update message.
- Behavior for mapping of PCE-initiated point-to-multipoint LSP with MVPN routing-instance and mapping of static (locally configured) point-to-multipoint LSP with MVPN instance is the same at user level.
- A flow specification ID can be associated with only one point-to-multipoint LSP. To associate the same RD and (S,G) to multiple point-to-multipoint LSPs, you can add multiple flow specifications with different IDs and same RD & (S,G).
- For PCEP-mapped dynamic (S,G), the threshold value is always the default value of 0.
- There is no limit on the number of flow specifications mapped to a single PCE-initiated point-to-multipoint LSP.
- The current implementation of this feature does not support:
 - Reporting of forwarding states that are associated with the point-to-multipoint LSP.
 - Inclusive provider tunnel dynamic configuration
 - Mapping for MVPN ingress replication tunnel
 - Programmable routing protocol process (prpd)
 - Reporting of CLI configured point-to-multipoint LSP which is mapped to MVPN multicast flows (S,G).

Release History Table

Release	Description
19.4R1	You can associate a single or range of MVPN multicast flows (S,G) to a dynamically created PCE-initiated point-to-multipoint label-switched path (LSP).

RELATED DOCUMENTATION

[pce | 3084](#)

[show path-computation-client lsp | 3559](#)

[show path-computation-client status | 3573](#)

[Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1734](#)

Configuring PCEP for MPLS SPRING-TE

IN THIS CHAPTER

- [Support of SPRING-TE for the Path Computation Element Protocol Overview | 1777](#)
- [Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs | 1782](#)
- [Static Segment Routing Label Switched Path | 1811](#)

Support of SPRING-TE for the Path Computation Element Protocol Overview

IN THIS SECTION

- [SPRING for Traffic Engineering | 1777](#)
- [Junos OS Implementation of PCEP for SPRING-TE LSPs | 1778](#)
- [Configuration of PCEP for SPRING-TE | 1780](#)
- [Limitations and Unsupported Features for PCEP SPRING-TE | 1781](#)

The traffic engineering (TE) capabilities of Source Packet Routing in Networking (SPRING) are supported in the Path Computation Element Protocol (PCEP) implementation of Junos OS. With this support, the advantages of SPRING are extended to the label-switched paths (LSPs) initiated by a Path Computation Element (PCE), augmenting the benefits of external path computing in an MPLS network. This topic describes the PCEP implementation for SPRING-TE LSPs.

SPRING for Traffic Engineering

Traditionally, RSVP has been used to address traffic engineering problems. RSVP-TE is an extension to RSVP that allows MPLS labels to be generated for prefixes, and at the same time uses the resource reservation capabilities to reserve specific LSPs through the network. Currently, the Junos OS

implementation of PCEP uses RSVP-TE. Starting in Junos OS Release 17.2, the PCEP implementation is extended to support SPRING-TE as well.

SPRING can operate over an IPv4 or IPv6 data plane, and supports equal-cost multipath (ECMP). With the IGP extensions built into it, SPRING integrates with the rich multiservice capabilities of MPLS, including Layer 3 VPN (L3VPN), Virtual Private Wire Service (VPWS), Virtual Private LAN Service (VPLS), and Ethernet VPN (EVPN).

Some of the high-level components of the SPRING-TE solution are:

- Use of an IGP for advertising link characteristics. This is similar to RSVP-TE.
- Use of CSPF on the ingress device or the PCE.
- Use of an IGP for advertising labels for links.

With SPRING, the ingress device constructs an LSP by stacking the labels of the links that it wants to traverse. The per-link IGP advertisement is combined with label stacking to create source routed LSPs on the ingress, so the transit devices are not aware of the end-to-end LSPs. As a result, there is no per-LSP signaling in SPRING-TE, and the ability to stack per-neighbor labels contributes to the control plane scaling property. This enables creation of LSPs between edge nodes without placing any per-LSP memory requirements on the transit devices. However, the label stacking feature of SPRING-TE results in the management of a large number of labels, which might be difficult for some platforms to support.

Junos OS Implementation of PCEP for SPRING-TE LSPs

IN THIS SECTION

- [SPRING-TE Module | 1779](#)
- [Traffic Engineering Database | 1779](#)
- [PCEP Interaction | 1779](#)

The Junos OS implementation of SPRING-TE LSPs for PCEP allows creation of tunnel routes for SPRING-TE LSPs, and using the tunnel routes for placing IP traffic or services. The SPRING-TE LSPs are created by the PCE using the adjacency and node segments.

Junos OS SPRING-TE LSPs cannot be longer than six hops. The PCE computes the path of the SPRING-TE LSP, and provisions the LSP on the Path Computation Client (PCC) using PCEP segment routing (SR) extensions. The PCEP SR extensions are parsed and a tunnel is created on the PCC. Like any other tunnel route, IP traffic and services can be resolved over the tunnel.

The Junos OS implementation of PCEP for SPRING-TE LSPs includes the following components:

SPRING-TE Module

SPRING-TE LSPs are created on the PCC, making the tunnel routes available in the inet.3 routing table. The PCC selects the outgoing interface based on the first network access identifier (NAI) in the source Explicit Route Object (S-ERO). If the PCC receives an S-ERO that does not have labels in it, it rejects the S-ERO. SPRING-TE creates tunnel routes that have their own preference value.

If the path has more than five labels, then the path is unviable. So, any S-ERO that carries more than six hops is rejected. When there are multiple LSPs to the same destination with the same metric, the PCC creates an equal-cost multipath (ECMP) route.

Junos OS supports S-EROs that contain the first hop as a strict hop; there is no support on the PCC to select the outgoing interface based on a loose hop node segment ID. However, the remaining hops can be loose. No specific processing is done for the S-EROs that are beyond the first hop, other than to simply use the label for next-hop creation.

Traffic Engineering Database

After a PCE has provisioned a SPRING-TE LSP for the PCC, a few things can change, such as:

- The label used by a path can change.
- The label used by a path can be withdrawn.
- One of the interfaces traversed by the LSP can go down, which is detected by monitoring IGP advertisements.

The PCC does not react to any of the above changes. The PCC waits for the PCE to process the event, and either re-program the LSP or bring down the LSP, as appropriate. However, when the directly attached interface of a SPRING-TE LSP goes down, that SPRING-TE LSP is not used for tunnel route creation.

PCEP Interaction

The Junos OS implementation of PCEP provides support to process received provisioning type S-ERO objects on the PCC. If the PCC loses connection with the PCE, the SPRING-TE LSPs remain up for 300 seconds, and then get purged out. If the PCE supplies parameters, such as setup priority, that do not apply to SR, the PCC ignores them.

The PCEP interaction with SPRING-TE LSPs is based on the Internet drafts—*draft-ietf-pce-lsp-setup-type-03.txt* and *draft-ietf-pce-segment-routing-06.txt*. The PCEP draft compliance is as follows:

- Support the PCE capability type, length, and value TLV. The maximum service identifier (SID) is set to 5 by default. The **max-sid-depth** statement can be used to control the maximum SID depth advertised.
- Path setup TLV is supported.
- ERO is supported, but loose hop expansion is not supported.
- Only IPv4 node IDs and adjacencies are supported. IPv6 or unnumbered adjacencies are not supported.
- Because the maximum SID depth is always set to five, the metric object is also set to five.

NOTE: The S-ERO objects that do not carry labels are rejected.

Configuration of PCEP for SPRING-TE

Enabling SPRING for PCEP

To enable SPRING for PCEP, you need to configure MPLS and SPRING-TE.

NOTE: PCEP for RSVP-TE cannot be disabled when PCEP for SPRING-TE is enabled.

To enable PCEP for SPRING-TE LSPs, the following configuration must be executed:

1. Enable external path computing capabilities for MPLS. This configuration is required for PCEP for RSVP-TE as well.

```
[edit protocols]
user@host# set mpls lsp-external-controller
```

2. Configure SPRING-TE with external path computing capabilities.

```
[edit protocols]
user@host# set spring-traffic-engineering lsp-external-controller pccd
```

3. Enable SPRING on a PCE.

```
[edit protocols]
user@host# set pcep pce pce1 spring-capability
```

Configuring Maximum SID Depth

The maximum SID depth value for SPRING-TE is set to five by default. The following configuration can be used to control the maximum SID depth that is advertised:

```
[edit protocols]
user@host# set pcep pce pce1 max-sid-depth 5
```

Configuring Preference Value for SPRING-TE

The default preference value of SPRING-TE is eight.

```
[edit protocols]
user@host# set spring-te preference preference-value
```

Configuring SPRING-TE Logging

The following configuration is used to control logging for SPRING-TE LSPs:

```
[edit protocols]
user@host# set spring-te traceoptions file file-name size file-size
user@host# set spring-te traceoptions flag [controller | state | route | general | interface | all]
```

Viewing SPRING-TE LSPs

The **show spring-te lsp** command displays the SPRING-TE LSPs on the device.

```
show spring-te lsp name lsp-name
```

```
show spring-te lsp extensive
```

Limitations and Unsupported Features for PCEP SPRING-TE

The support of SPRING for PCEP does not add any additional performance burden on the system; however, the PCEP SPRING-TE implementation has the following limitations:

- A SPRING-TE LSP is not locally protected on the PCC. When the LSP is over six hops, no other service is provided on the LSP other than to carry plain IP.
- Graceful Routing Engine switchover (GRES) and unified ISSU in-service software upgrade are not supported.
- Nonstop active routing (NSR) is not supported.
- IPv6 is not supported.
- Logical systems are not supported.

RELATED DOCUMENTATION

[Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs](#) | 1782

Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs

IN THIS SECTION

- [Requirements | 1782](#)
- [Overview | 1782](#)
- [Configuration | 1783](#)
- [Verification | 1791](#)

This example shows how to configure Path Computation Element Protocol (PCEP) for traffic engineered label-switched paths (LSPs) of Source Packet Routing in Networking (SPRING). The advantages of SPRING are leveraged with the benefits of external path computing for efficient traffic engineering.

Requirements

This example uses the following hardware and software components:

- Four MX Series 5G Universal Routing Platforms, where the ingress router is the Path Computation Client (PCC).
- A TCP connection to an external stateful Path Computation Element (PCE) from the PCC.
- Junos OS Release 17.2 or later running on the PCC.

Before you begin:

- Configure the device interfaces.
- Configure MPLS.
- Configure IS-IS.

Overview

Starting in Junos OS Release 17.2, the traffic engineering capabilities of Source Packet Routing in Networking (SPRING) are supported in PCEP sessions for the LSPs initiated by a PCE. Tunnel routes are created in the inet.3 routing table of the PCC corresponding to the SPRING-TE LSPs. Similar to any other tunnel route, the SPRING-TE tunnel routes can be used for resolving indirect next hops for plain IP and service traffic. The SPRING-TE LSPs are created by the PCE for the adjacency and node segments.

The SPRING-TE LSPs can be a maximum of six hops long. The PCE computes the path of the SPRING-TE LSP, and provisions the LSP on the Path Computation Client (PCC) using PCEP segment routing (SR) extensions. Tunnel routes are created in the inet.3 routing table of the PCC corresponding to the SPRING-TE LSPs. Similar to any other tunnel route, the SPRING-TE tunnel routes can be used for resolving indirect next hops for plain IP and service traffic.


To configure SPRING-TE for PCEP:

- Enable external path computing for MPLS and SPRING-TE at the **[edit protocols]** hierarchy level.
- Enable spring capability for the PCE at the **[edit protocols pcep pce pce]** hierarchy level.

Topology

Figure 143 on page 1783 illustrates a sample network topology that has PCE as the external path computing entity, and the PCC as the ingress router that connects to the PCE. Routers R1, R2, and R3 are the other routers in the network. The PCC is enabled with SPRING-TE capability for the PCEP sessions. A static route is configured on the PCC to Router R3 to verify the use of SPRING-TE tunnel routes when routing traffic for the static route.

Figure 143: PCEP for SPRING-TE LSPs



Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

PCC


```

set interfaces ge-0/0/5 unit 0 family inet address 172.16.0.1/16
set interfaces ge-0/0/5 unit 0 family iso
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.4/8 primary
set interfaces lo0 unit 0 family iso address 49.0011.0110.0000.0101.00
set interfaces lo0 unit 0 family mpls
set routing-options static route 100.1.1.1/32 next-hop 127.0.0.3
set routing-options router-id 127.0.0.4
set routing-options autonomous-system 11
set protocols rsvp interface fxp0.0 disable
set protocols rsvp interface all
set protocols mpls lsp-external-controller pccd
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis source-packet-routing srgb start-label 800000
set protocols isis source-packet-routing srgb index-range 4000
set protocols isis source-packet-routing node-segment ipv4-index 101
set protocols isis source-packet-routing node-segment ipv6-index 11
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive
set protocols spring-traffic-engineering lsp-external-controller pccd
set protocols pcep pce pce1 local-address 127.0.0.4
set protocols pcep pce pce1 destination-ipv4-address 10.102.180.232
set protocols pcep pce pce1 destination-port 4189
set protocols pcep pce pce1 pce-type active
set protocols pcep pce pce1 pce-type stateful
set protocols pcep pce pce1 lsp-provisioning
set protocols pcep pce pce1 spring-capability

```

Router R1

```

set interfaces ge-0/0/5 unit 0 family inet address 172.16.0.2/16
set interfaces ge-0/0/5 unit 0 family iso
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/1/2 unit 0 family inet address 172.17.0.1/16
set interfaces ge-0/1/2 unit 0 family iso

```



```

set interfaces ge-0/1/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/8 primary
set interfaces lo0 unit 0 family iso address 49.0011.0110.0000.0102.00
set interfaces lo0 unit 0 family mpls
set routing-options router-id 127.0.0.1
set routing-options autonomous-system 11
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis source-packet-routing srgb start-label 800000
set protocols isis source-packet-routing srgb index-range 4000
set protocols isis source-packet-routing node-segment ipv4-index 102
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive

```

Router R2

```

set interfaces ge-0/1/2 unit 0 family inet address 172.17.0.2/16
set interfaces ge-0/1/2 unit 0 family iso
set interfaces ge-0/1/2 unit 0 family mpls
set interfaces ge-0/1/8 unit 0 family inet address 172.18.0.1/16
set interfaces ge-0/1/8 unit 0 family iso
set interfaces ge-0/1/8 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.2/8
set interfaces lo0 unit 0 family iso address 49.0011.0110.0000.0105.00
set interfaces lo0 unit 0 family mpls
set routing-options router-id 127.0.0.2
set routing-options autonomous-system 11
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis source-packet-routing srgb start-label 800000
set protocols isis source-packet-routing srgb index-range 4000
set protocols isis source-packet-routing node-segment ipv4-index 105

```



```

set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface all level 1 disable
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive

```

Router R3

```

set interfaces ge-0/1/8 unit 0 family inet address 172.18.0.2/16
set interfaces ge-0/1/8 unit 0 family iso
set interfaces ge-0/1/8 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.3/8 primary
set interfaces lo0 unit 0 family iso address 49.0011.0110.0000.0103.00
set interfaces lo0 unit 0 family mpls
set routing-options static route 100.1.1.1/32 receive
set routing-options router-id 127.0.0.3
set routing-options autonomous-system 11
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis source-packet-routing srgb start-label 800000
set protocols isis source-packet-routing srgb index-range 4000
set protocols isis source-packet-routing node-segment ipv4-index 103
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive

```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure PCC:

1. Configure the interfaces of the PCC.


```
[edit interfaces]
user@PCC# set ge-0/0/5 unit 0 family inet address 172.16.0.1/16
user@PCC# set ge-0/0/5 unit 0 family iso
user@PCC# set ge-0/0/5 unit 0 family mpls
user@PCC# set lo0 unit 0 family inet address 127.0.0.4/8 primary
user@PCC# set lo0 unit 0 family iso address 49.0011.0110.0000.0101.00
user@PCC# set lo0 unit 0 family mpls
```

2. Configure the router ID and assign an autonomous system number for the PCC.

```
[edit routing-options]
user@PCC# set router-id 127.0.0.4
user@PCC# set autonomous-system 11
```

3. Configure a static route from the PCC to Router R3.

The static route is created for verification purpose only and does not affect the feature functionality.

```
[edit routing-options]
user@PCC# set static route 100.1.1.1/32 next-hop 127.0.0.3
```

4. Configure RSVP on all the interfaces of the PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set rsvp interface fxp0.0 disable
user@PCC# set rsvp interface all
```

5. Configure MPLS on all the interfaces of the PCC, excluding the management interface.

```
[edit protocols]
user@PCC# set mpls interface all
user@PCC# set mpls interface fxp0.0 disable
```

6. Enable external path computing capability for the PCC.

```
[edit protocols]
user@PCC# set mpls lsp-external-controller pccd
```


7. Configure IS-IS level 2 on all the interfaces of the PCC, excluding the management and loopback interfaces.

```
[edit protocols]
user@PCC# set isis level 1 disable
user@PCC# set isis level 2 wide-metrics-only
user@PCC# set isis interface all point-to-point
user@PCC# set isis interface all level 2 metric 10
user@PCC# set isis interface fxp0.0 disable
user@PCC# set isis interface lo0.0 passive
```

8. Configure Segment routing global block (SRGB) attributes for SPRING.

```
[edit protocols]
user@PCC# set isis source-packet-routing srgb start-label 800000
user@PCC# set isis source-packet-routing srgb index-range 4000
user@PCC# set isis source-packet-routing node-segment ipv4-index 101
user@PCC# set isis source-packet-routing node-segment ipv6-index 11
```

9. Configure SPRING with external path computing capability.

```
[edit protocols]
user@PCC# set spring-traffic-engineering lsp-external-controller pccd
```

10. Configure the PCE parameters and enable provisioning of LSP by the PCE and SPRING capability.

```
[edit protocols]
user@PCC# set pcep pce pce1 local-address 127.0.0.4
user@PCC# set pcep pce pce1 destination-ipv4-address 10.102.180.232
user@PCC# set pcep pce pce1 destination-port 4189
user@PCC# set pcep pce pce1 pce-type active
user@PCC# set pcep pce pce1 pce-type stateful
user@PCC# set pcep pce pce1 lsp-provisioning
user@PCC# set pcep pce pce1 spring-capability
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.


```

user@PCC# show interfaces
ge-0/0/5 {
  unit 0 {
    family inet {
      address 172.16.0.1/16;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.4/8 {
        primary;
      }
    }
    family iso {
      address 49.0011.0110.0000.0101.00;
    }
    family mpls;
  }
}

```

```

user@PCC# show routing-options
static {
  route 100.1.1.1/32 next-hop 127.0.0.3;
}
router-id 127.0.0.4;
autonomous-system 11;

```

```

user@PCC# show protocols
rsvp {
  interface fxp0.0 {
    disable;
  }
  interface all;
}
mpls {
  lsp-external-controller pccd;
  interface all;
  interface fxp0.0 {
    disable;
  }
}

```



```

    }
}
isis {
    source-packet-routing {
        srgb start-label 800000 index-range 4000;
        node-segment {
            ipv4-index 101;
            ipv6-index 11;
        }
    }
    level 1 disable;
    level 2 wide-metrics-only;
    interface all {
        point-to-point;
        level 2 metric 10;
    }
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        passive;
    }
}
spring-traffic-engineering {
    lsp-external-controller pccd;
}
pcep {
    pce pce1 {
        local-address 127.0.0.4;
        destination-ipv4-address 10.102.180.232;
        destination-port 4189;
        pce-type active stateful;
        lsp-provisioning;
        spring-capability;
    }
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verifying IS-IS Adjacency and Labels | 1791](#)
- [Verifying the Traffic Engineering Database | 1800](#)
- [Verifying SPRING-TE LSPs | 1804](#)
- [Verifying Tunnel Route Creation | 1806](#)
- [Verifying Forwarding Table Entries | 1808](#)
- [Verifying Use of Tunnel Routes for Static Route Forwarding | 1810](#)

Confirm that the configuration is working properly.

Verifying IS-IS Adjacency and Labels

Purpose

Verify the IS-IS adjacency and label assignments for interfaces and nodes on the PCC.

Action

From operational mode, run the **show isis adjacency extensive**, **show isis database extensive**, and **show isis overview** commands.

```
user@PCC> show isis adjacency extensive
```

```
R1
Interface: ge-0/0/5.0, Level: 2, State: Up, Expires in 25 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:37:15 ago
Circuit type: 2, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 172.16.0.2
Level 2 IPv4 Adj-SID: 16
Transition log:
When                State      Event      Down reason
Wed Apr  5 02:42:48  Up        Seenself

PCE
Interface: gre.0, Level: 2, State: Up, Expires in 25 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:27:00 ago
```



```

Circuit type: 2, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes, Adjacency advertisement: Advertise
IP addresses: 11.105.199.2
Level 2
Transition log:

```

When	State	Event	Down reason
Wed Apr 5 02:53:03	Up	Seenself	

user@PCC> **show isis database extensive**

```

IS-IS level 1 link-state database:

IS-IS level 2 link-state database:

PCC.00-00 Sequence: 0x2a6, Checksum: 0x1a4f, Lifetime: 1150 secs
  IPV4 Index: 101
  Node Segment Blocks Advertised:
    Start Index : 0, Size : 4000, Label-Range: [ 800000, 803999 ]
    IS neighbor: R1.00                               Metric: 10
      Two-way fragment: R1.00-00, Two-way first fragment: R1.00-00
    IS neighbor: PCE.00                               Metric: 16777215
    IP prefix: 127.0.0.4/8                             Metric: 0 Internal Up
    IP prefix: 11.101.102.0/30                         Metric: 10 Internal Up
    IP prefix: 11.105.199.0/30                         Metric: 16777215 Internal Up

Header: LSP ID: PCC.00-00, Length: 243 bytes
  Allocated length: 1492 bytes, Router ID: 127.0.0.4
  Remaining lifetime: 1150 secs, Level: 2, Interface: 0
  Estimated free bytes: 1084, Actual free bytes: 1249
  Aging timer expires in: 1150 secs
  Protocols: IP, IPv6

Packet: LSP ID: PCC.00-00, Length: 243 bytes, Lifetime : 1198 secs
  Checksum: 0x1a4f, Sequence: 0x2a6, Attributes: 0x3 L1 L2
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.0011 (3)
  LSP Buffer Size: 1492
  Speaks: IP

```



```

Speaks: IPV6
IP router id: 127.0.0.4
IP address: 127.0.0.4
Hostname: PCC
IS extended neighbor: R1.00, Metric: default 10
  IP address: 172.16.0.1
  Neighbor's IP address: 172.16.0.2
  Local interface index: 334, Remote interface index: 333
  Current reservable bandwidth:
    Priority 0 : 10Mbps
    Priority 1 : 10Mbps
    Priority 2 : 10Mbps
    Priority 3 : 10Mbps
    Priority 4 : 10Mbps
    Priority 5 : 10Mbps
    Priority 6 : 10Mbps
    Priority 7 : 10Mbps
  Maximum reservable bandwidth: 10Mbps
  Maximum bandwidth: 10Mbps
  Administrative groups: 0 none
  P2P IPV4 Adj-SID - Flags:0x30(F:0,B:0,V:1,L:1,S:0), Weight:0, Label: 16
IS extended neighbor: PCE.00, Metric: default 16777215
  IP address: 11.105.199.1
  Neighbor's IP address: 11.105.199.2
  Local interface index: 329, Remote interface index: 329
IP extended prefix: 11.101.102.0/30 metric 10 up
IP extended prefix: 11.105.199.0/30 metric 16777215 up
IP extended prefix: 127.0.0.4/8 metric 0 up
  8 bytes of subtlvs
  Node SID, Flags: 0x40(R:0,N:1,P:0,E:0,V:0,L:0), Algo: SPF(0), Value: 101
Router Capability: Router ID 127.0.0.4, Flags: 0x00
  SPRING Capability - Flags: 0xc0(I:1,V:1), Range: 4000, SID-Label: 800000
  SPRING Algorithm - Algo: 0
No queued transmissions

R1.00-00 Sequence: 0x297, Checksum: 0x1615, Lifetime: 839 secs
IPV4 Index: 102
Node Segment Blocks Advertised:
  Start Index : 0, Size : 4000, Label-Range: [ 800000, 803999 ]
IS neighbor: PCC.00 Metric: 10
  Two-way fragment: PCC.00-00, Two-way first fragment: PCC.00-00
IS neighbor: R2.00 Metric: 10
  Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
IP prefix: 127.0.0.1/8 Metric: 0 Internal Up

```



```

IP prefix: 11.101.102.0/30          Metric:      10 Internal Up
IP prefix: 11.102.105.0/30          Metric:      10 Internal Up

```

```

Header: LSP ID: R1.00-00, Length: 302 bytes
  Allocated length: 302 bytes, Router ID: 127.0.0.1
  Remaining lifetime: 839 secs, Level: 2, Interface: 334
  Estimated free bytes: 0, Actual free bytes: 0
  Aging timer expires in: 839 secs
  Protocols: IP, IPv6

```

```

Packet: LSP ID: R1.00-00, Length: 302 bytes, Lifetime : 1196 secs
  Checksum: 0x1615, Sequence: 0x297, Attributes: 0x3 L1 L2
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

TLVs:

```

  Area address: 49.0011 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 127.0.0.1
  IP address: 127.0.0.1
  Hostname: R1
  IP extended prefix: 127.0.0.1/8 metric 0 up
    8 bytes of subtlvs
    Node SID, Flags: 0x40(R:0,N:1,P:0,E:0,V:0,L:0), Algo: SPF(0), Value: 102
  IP extended prefix: 11.101.102.0/30 metric 10 up
  IP extended prefix: 11.102.105.0/30 metric 10 up
  Router Capability: Router ID 127.0.0.1, Flags: 0x00
    SPRING Capability - Flags: 0xc0(I:1,V:1), Range: 4000, SID-Label: 800000
    SPRING Algorithm - Algo: 0
  IS extended neighbor: R2.00, Metric: default 10
    IP address: 172.17.0.1
    Neighbor's IP address: 172.17.0.2
    Local interface index: 334, Remote interface index: 333
  Current reservable bandwidth:
    Priority 0 : 10Mbps
    Priority 1 : 10Mbps
    Priority 2 : 10Mbps
    Priority 3 : 10Mbps
    Priority 4 : 10Mbps
    Priority 5 : 10Mbps
    Priority 6 : 10Mbps
    Priority 7 : 10Mbps

```



```

Maximum reservable bandwidth: 10Mbps
Maximum bandwidth: 10Mbps
Administrative groups: 0 none
P2P IPV4 Adj-SID - Flags:0x30(F:0,B:0,V:1,L:1,S:0), Weight:0, Label: 17
IS extended neighbor: PCC.00, Metric: default 10
IP address: 172.16.0.2
Neighbor's IP address: 172.16.0.1
Local interface index: 333, Remote interface index: 334
Current reservable bandwidth:
    Priority 0 : 10Mbps
    Priority 1 : 10Mbps
    Priority 2 : 10Mbps
    Priority 3 : 10Mbps
    Priority 4 : 10Mbps
    Priority 5 : 10Mbps
    Priority 6 : 10Mbps
    Priority 7 : 10Mbps
Maximum reservable bandwidth: 10Mbps
Maximum bandwidth: 10Mbps
Administrative groups: 0 none
P2P IPV4 Adj-SID - Flags:0x30(F:0,B:0,V:1,L:1,S:0), Weight:0, Label: 16
No queued transmissions

R3.00-00 Sequence: 0x95, Checksum: 0xd459, Lifetime: 895 secs
IPV4 Index: 103
Node Segment Blocks Advertised:
    Start Index : 0, Size : 4000, Label-Range: [ 800000, 803999 ]
IS neighbor: R2.00 Metric: 10
    Two-way fragment: R2.00-00, Two-way first fragment: R2.00-00
IP prefix: 127.0.0.3/8 Metric: 0 Internal Up
IP prefix: 11.102.1.0/24 Metric: 10 Internal Up
IP prefix: 11.103.107.0/30 Metric: 10 Internal Up

Header: LSP ID: R3.00-00, Length: 209 bytes
    Allocated length: 284 bytes, Router ID: 127.0.0.3
    Remaining lifetime: 895 secs, Level: 2, Interface: 334
    Estimated free bytes: 75, Actual free bytes: 75
    Aging timer expires in: 895 secs
    Protocols: IP, IPv6

Packet: LSP ID: R3.00-00, Length: 209 bytes, Lifetime : 1192 secs
    Checksum: 0xd459, Sequence: 0x95, Attributes: 0x3 L1 L2
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

```


TLVs:

```

Area address: 49.0011 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPV6
IP router id: 127.0.0.3
IP address: 127.0.0.3
Hostname: R3
IS extended neighbor: R2.00, Metric: default 10
  IP address: 172.18.0.2
  Neighbor's IP address: 172.18.0.1
  Local interface index: 336, Remote interface index: 334
  Current reservable bandwidth:
    Priority 0 : 10Mbps
    Priority 1 : 10Mbps
    Priority 2 : 10Mbps
    Priority 3 : 10Mbps
    Priority 4 : 10Mbps
    Priority 5 : 10Mbps
    Priority 6 : 10Mbps
    Priority 7 : 10Mbps
  Maximum reservable bandwidth: 10Mbps
  Maximum bandwidth: 10Mbps
  Administrative groups: 0 none
  P2P IPV4 Adj-SID - Flags:0x30(F:0,B:0,V:1,L:1,S:0), Weight:0, Label: 16
IP extended prefix: 127.0.0.3/8 metric 0 up
  8 bytes of subtlvs
  Node SID, Flags: 0x40(R:0,N:1,P:0,E:0,V:0,L:0), Algo: SPF(0), Value: 103
IP extended prefix: 11.103.107.0/30 metric 10 up
IP extended prefix: 11.102.1.0/24 metric 10 up
Router Capability: Router ID 127.0.0.3, Flags: 0x00
  SPRING Capability - Flags: 0xc0(I:1,V:1), Range: 4000, SID-Label: 800000
  SPRING Algorithm - Algo: 0
No queued transmissions

```

R2.00-00 Sequence: 0x2aa, Checksum: 0xf8f4, Lifetime: 1067 secs

IPV4 Index: 105

Node Segment Blocks Advertised:

Start Index : 0, Size : 4000, Label-Range: [800000, 803999]

IS neighbor: R1.00 Metric: 10

Two-way fragment: R1.00-00, Two-way first fragment: R1.00-00

IS neighbor: R3.00 Metric: 10

Two-way fragment: R3.00-00, Two-way first fragment: R3.00-00

IP prefix: 127.0.0.2/8	Metric:	0 Internal Up
IP prefix: 11.102.105.0/30	Metric:	10 Internal Up
IP prefix: 11.103.107.0/30	Metric:	10 Internal Up

Header: LSP ID: R2.00-00, Length: 302 bytes
 Allocated length: 302 bytes, Router ID: 127.0.0.2
 Remaining lifetime: 1067 secs, Level: 2, Interface: 334
 Estimated free bytes: 0, Actual free bytes: 0
 Aging timer expires in: 1067 secs
 Protocols: IP, IPv6

Packet: LSP ID: R2.00-00, Length: 302 bytes, Lifetime : 1194 secs
 Checksum: 0xf8f4, Sequence: 0x2aa, Attributes: 0x3 L1 L2
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0011 (3)
 LSP Buffer Size: 1492
 Speaks: IP
 Speaks: IPV6
 IP router id: 127.0.0.2
 IP address: 127.0.0.2
 Hostname: R2
 IP extended prefix: 127.0.0.2/8 metric 0 up
 8 bytes of subtlvs
 Node SID, Flags: 0x40(R:0,N:1,P:0,E:0,V:0,L:0), Algo: SPF(0), Value: 105
 IP extended prefix: 11.102.105.0/30 metric 10 up
 IP extended prefix: 11.103.107.0/30 metric 10 up
 Router Capability: Router ID 127.0.0.2, Flags: 0x00
 SPRING Capability - Flags: 0xc0(I:1,V:1), Range: 4000, SID-Label: 800000
 SPRING Algorithm - Algo: 0
 IS extended neighbor: R3.00, Metric: default 10
 IP address: 172.18.0.1
 Neighbor's IP address: 172.18.0.2
 Local interface index: 334, Remote interface index: 336
 Current reservable bandwidth:
 Priority 0 : 10Mbps
 Priority 1 : 10Mbps
 Priority 2 : 10Mbps
 Priority 3 : 10Mbps
 Priority 4 : 10Mbps
 Priority 5 : 10Mbps
 Priority 6 : 10Mbps


```

    Priority 7 : 10Mbps
Maximum reservable bandwidth: 10Mbps
Maximum bandwidth: 10Mbps
Administrative groups: 0 none
P2P IPV4 Adj-SID - Flags:0x30(F:0,B:0,V:1,L:1,S:0), Weight:0, Label: 16
IS extended neighbor: R1.00, Metric: default 10
IP address: 172.17.0.2
Neighbor's IP address: 172.17.0.1
Local interface index: 333, Remote interface index: 334
Current reservable bandwidth:
    Priority 0 : 10Mbps
    Priority 1 : 10Mbps
    Priority 2 : 10Mbps
    Priority 3 : 10Mbps
    Priority 4 : 10Mbps
    Priority 5 : 10Mbps
    Priority 6 : 10Mbps
    Priority 7 : 10Mbps
Maximum reservable bandwidth: 10Mbps
Maximum bandwidth: 10Mbps
Administrative groups: 0 none
P2P IPV4 Adj-SID - Flags:0x30(F:0,B:0,V:1,L:1,S:0), Weight:0, Label: 17
No queued transmissions

PCE.00-00 Sequence: 0x277, Checksum: 0x64a5, Lifetime: 533 secs
IS neighbor: PCC.00 Metric: 16777215
IP prefix: 11.0.0.199/32 Metric: 0 Internal Up
IP prefix: 11.105.199.0/30 Metric: 16777215 Internal Up

Header: LSP ID: PCE.00-00, Length: 120 bytes
Allocated length: 284 bytes, Router ID: 11.0.0.199
Remaining lifetime: 533 secs, Level: 2, Interface: 329
Estimated free bytes: 164, Actual free bytes: 164
Aging timer expires in: 533 secs
Protocols: IP, IPv6

Packet: LSP ID: PCE.00-00, Length: 120 bytes, Lifetime : 1196 secs
Checksum: 0x64a5, Sequence: 0x277, Attributes: 0x3 L1 L2
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

TLVs:
Area address: 11.0007 (3)
LSP Buffer Size: 1492

```



```

Speaks: IP
Speaks: IPV6
IP router id: 11.0.0.199
IP address: 11.0.0.199
Hostname: PCE
Router Capability: Router ID 11.0.0.199, Flags: 0x00
IP extended prefix: 11.105.199.0/30 metric 16777215 up
IP extended prefix: 11.0.0.199/32 metric 0 up
IS extended neighbor: PCC.00, Metric: default 16777215
  IP address: 11.105.199.2
  Neighbor's IP address: 11.105.199.1
  Local interface index: 329, Remote interface index: 329
No queued transmissions

```

user@PCC> **show isis overview**

```

Instance: master
  Router ID: 127.0.0.4
  Hostname: PCC
  Sysid: 0110.0000.0101
  Areaid: 49.0011
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled, SPRING based MPLS is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
  Layer2-map: Disabled
  Source Packet Routing (SPRING): Enabled
  SRGB Config Range:
    SRGB Start-Label : 800000, SRGB Index-Range : 4000
  SRGB Block Allocation: Success
    SRGB Start Index : 800000, SRGB Size : 4000, Label-Range: [ 800000, 803999 ]
  Node Segments: Enabled
    Ipv4 Index : 101, Ipv6 Index : 11
Level 1
  Internal route preference: 15
  External route preference: 160
  Prefix export count: 0
  Wide metrics are enabled, Narrow metrics are enabled

```



```

Source Packet Routing is enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export count: 0
  Wide metrics are enabled
  Source Packet Routing is enabled

```

Meaning

The IS-IS adjacency between the PCC and PCE and the PCC and Router R1 is up and operational. The output also displays the label assignments for the adjacent and node segments.

Verifying the Traffic Engineering Database

Purpose

Verify the traffic engineering database entries on the PCC.

Action

From operational mode, run the **show ted database extensive** command.

user@PCC# **show ted database extensive**

```

TED database: 5 ISIS nodes 5 INET nodes
NodeID: PCC.00(127.0.0.4)
  Type: Rtr, Age: 403 secs, LinkIn: 1, LinkOut: 1
  Protocol: IS-IS(2)
  127.0.0.4
    To: R1.00(127.0.0.1), Local: 172.16.0.1, Remote: 172.16.0.2
    Local interface index: 334, Remote interface index: 333
    Color: 0 none
    Metric: 10
    IGP metric: 10
    Static BW: 10Mbps
    Reservable BW: 10Mbps
    Available BW [priority] bps:
      [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
      [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
        [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps

```



```

P2P Adjacency-SID:
  IPV4, SID: 16, Flags: 0x30, Weight: 0
Prefixes:
  127.0.0.4/8
  Metric: 0, Flags: 0x00
  Prefix-SID:
    SID: 101, Flags: 0x40, Algo: 0
SPRING-Capabilities:
  SRGB block [Start: 800000, Range: 4000, Flags: 0xc0]
SPRING-Algorithms:
  Algo: 0
NodeID: R1.00(127.0.0.1)
Type: Rtr, Age: 712 secs, LinkIn: 2, LinkOut: 2
Protocol: IS-IS(2)
127.0.0.1
To: PCC.00(127.0.0.4), Local: 172.16.0.2, Remote: 172.16.0.1
  Local interface index: 333, Remote interface index: 334
  Color: 0 none
  Metric: 10
  IGP metric: 10
  Static BW: 10Mbps
  Reservable BW: 10Mbps
  Available BW [priority] bps:
    [0] 10Mbps    [1] 10Mbps    [2] 10Mbps    [3] 10Mbps
    [4] 10Mbps    [5] 10Mbps    [6] 10Mbps    [7] 10Mbps
  Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
      [0] 10Mbps    [1] 10Mbps    [2] 10Mbps    [3] 10Mbps
      [4] 10Mbps    [5] 10Mbps    [6] 10Mbps    [7] 10Mbps
P2P Adjacency-SID:
  IPV4, SID: 16, Flags: 0x30, Weight: 0
To: R2.00(127.0.0.2), Local: 172.17.0.1, Remote: 172.17.0.2
  Local interface index: 334, Remote interface index: 333
  Color: 0 none
  Metric: 10
  IGP metric: 10
  Static BW: 10Mbps
  Reservable BW: 10Mbps
  Available BW [priority] bps:
    [0] 10Mbps    [1] 10Mbps    [2] 10Mbps    [3] 10Mbps
    [4] 10Mbps    [5] 10Mbps    [6] 10Mbps    [7] 10Mbps
  Interface Switching Capability Descriptor(1):

```



```

    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
        [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    P2P Adjacency-SID:
        IPV4, SID: 17, Flags: 0x30, Weight: 0
Prefixes:
    127.0.0.1/8
    Metric: 0, Flags: 0x00
    Prefix-SID:
        SID: 102, Flags: 0x40, Algo: 0
SPRING-Capabilities:
    SRGB block [Start: 800000, Range: 4000, Flags: 0xc0]
SPRING-Algorithms:
    Algo: 0
NodeID: R3.00(127.0.0.3)
Type: Rtr, Age: 435 secs, LinkIn: 1, LinkOut: 1
Protocol: IS-IS(2)
127.0.0.3
To: R2.00(127.0.0.2), Local: 172.18.0.2, Remote: 172.18.0.1
    Local interface index: 336, Remote interface index: 334
    Color: 0 none
    Metric: 10
    IGP metric: 10
    Static BW: 10Mbps
    Reservable BW: 10Mbps
    Available BW [priority] bps:
        [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
        [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
            [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
            [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    P2P Adjacency-SID:
        IPV4, SID: 16, Flags: 0x30, Weight: 0
Prefixes:
    127.0.0.3/8
    Metric: 0, Flags: 0x00
    Prefix-SID:
        SID: 103, Flags: 0x40, Algo: 0
SPRING-Capabilities:

```



```

    SRGB block [Start: 800000, Range: 4000, Flags: 0xc0]
    SPRING-Algorithms:
        Algo: 0
NodeID: R2.00(127.0.0.2)
Type: Rtr, Age: 456 secs, LinkIn: 2, LinkOut: 2
Protocol: IS-IS(2)
127.0.0.2
To: R1.00(127.0.0.1), Local: 172.17.0.2, Remote: 172.17.0.1
    Local interface index: 333, Remote interface index: 334
    Color: 0 none
    Metric: 10
    IGP metric: 10
    Static BW: 10Mbps
    Reservable BW: 10Mbps
    Available BW [priority] bps:
        [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
        [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
            [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
            [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    P2P Adjacency-SID:
        IPV4, SID: 17, Flags: 0x30, Weight: 0
To: R3.00(127.0.0.3), Local: 172.18.0.1, Remote: 172.18.0.2
    Local interface index: 334, Remote interface index: 336
    Color: 0 none
    Metric: 10
    IGP metric: 10
    Static BW: 10Mbps
    Reservable BW: 10Mbps
    Available BW [priority] bps:
        [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
        [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
            [0] 10Mbps      [1] 10Mbps      [2] 10Mbps      [3] 10Mbps
            [4] 10Mbps      [5] 10Mbps      [6] 10Mbps      [7] 10Mbps
    P2P Adjacency-SID:
        IPV4, SID: 16, Flags: 0x30, Weight: 0
Prefixes:

```



```

127.0.0.2/8
Metric: 0, Flags: 0x00
Prefix-SID:
  SID: 105, Flags: 0x40, Algo: 0
SPRING-Capabilities:
  SRGB block [Start: 800000, Range: 4000, Flags: 0xc0]
SPRING-Algorithms:
  Algo: 0
NodeID: PCE.00(11.0.0.199)
Type: Rtr, Age: 267 secs, LinkIn: 0, LinkOut: 0
Protocol: IS-IS(2)
11.0.0.199

```

Meaning

The traffic engineering database includes entries advertised from Routers R1, R2, and R3, which the PCE uses for external path computing for the PCC.

Verifying *SPRING-TE LSPs*

Purpose

Verify the creation of *SPRING-TE LSP* on the PCC.

Action

From operational mode, run the **show path-computation-client lsp** and **show spring-traffic-engineering lsp detail** commands.

```
user@PCC> show path-computation-client lsp
```

Name	Path-Setup-Type	Status	PLSP-Id	LSP-Type
Controller		Template		
adj_sid_lsp		(Up)	3	ext-provided
pcel	spring-te			
node_sid_lsp		(Up)	5	ext-provided
pcel	spring-te			

```
user@PCC> show spring-traffic-engineering lsp detail
```

```

Name: adj_sid_lsp
To: 127.0.0.3

```



```

State: Up, Outgoing interface: ge-0/0/5.0
SR-ERO hop count: 3
  Hop 1 (Strict):
    NAI: IPv4 Adjacency ID, 172.16.0.1 -> 172.16.0.2
    SID type: 20-bit label, Value: 16
  Hop 2 (Strict):
    NAI: IPv4 Adjacency ID, 172.17.0.1 -> 172.17.0.2
    SID type: 20-bit label, Value: 17
  Hop 3 (Strict):
    NAI: IPv4 Adjacency ID, 172.18.0.1 -> 172.18.0.2
    SID type: 20-bit label, Value: 16

```

```

Name: node_sid_lsp
To: 127.0.0.3
State: Up, Outgoing interface: ge-0/0/5.0
SR-ERO hop count: 3
  Hop 1 (Strict):
    NAI: IPv4 Adjacency ID, 172.16.0.1 -> 172.16.0.2
    SID type: 20-bit label, Value: 16
  Hop 2 (Strict):
    NAI: IPv4 Node ID, Node address: 127.0.0.1
    SID type: 20-bit label, Value: 800105
  Hop 3 (Strict):
    NAI: IPv4 Node ID, Node address: 127.0.0.2
    SID type: 20-bit label, Value: 800103

```

```
Total displayed LSPs: 2 (Up: 2, Down: 0)
```

user@PCC> **show route protocol spring-te**

```

inet.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)

inet.3: 3 destinations, 4 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.3/8      *[SPRING-TE/8] 00:23:32, metric 0
                  to 172.16.0.2 via ge-0/0/5.0, Push 16, Push 17(top)
                  > to 172.16.0.2 via ge-0/0/5.0, Push 800103, Push 800105(top)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)

```



```
inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning

The outputs show that two SPRING-TE LSPs—**adj_sid_lsp** and **node_sid_lsp**—have been created by the PCE for the adjacency and node segments, respectively.

Verifying Tunnel Route Creation

Purpose

Verify the tunnel routes created for the SPRING-TE LSPs that are included in the inet.3 routing table on the PCC.

Action

From operation mode, run the **show route table inet.3 extensive** command.

```
user@PCC> show route table inet.3 extensive
```

```
inet.3: 3 destinations, 4 routes (3 active, 0 holddown, 0 hidden)
127.0.0.1/8 (1 entry, 1 announced)
  *L-ISIS Preference: 14
    Level: 2
    Next hop type: Router, Next hop index: 581
    Address: 0xb7a23b0
    Next-hop reference count: 13
    Next hop: 172.16.0.2 via ge-0/0/5.0, selected
    Session Id: 0xl72
    State: Active Int
    Local AS: 11
    Age: 45:51 Metric: 10
    Validation State: unverified
    ORR Generation-ID: 0
    Task: IS-IS
    Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 3
    AS path: I

127.0.0.3/8 (2 entries, 1 announced)
  *SPRING-TE Preference: 8
    Next hop type: Router, Next hop index: 0
    Address: 0xb61c190
    Next-hop reference count: 7
    Next hop: 172.16.0.2 via ge-0/0/5.0 weight 0x1
    Label operation: Push 16, Push 17(top)
```



```

Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 16: None; Label 17: None;
Label element ptr: 0xb7a2a60
Label parent element ptr: 0x0
Label element references: 5
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 172.16.0.2 via ge-0/0/5.0 weight 0x1, selected
Label operation: Push 800103, Push 800105(top)
Label TTL action: prop-ttl, prop-ttl(top)
Load balance label: Label 800103: None; Label 800105: None;
Label element ptr: 0xb7a2c40
Label parent element ptr: 0x0
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
State: Active Int
Local AS:      11
Age: 9:44      Metric: 0
Validation State: unverified
Task: SPRING-TE
Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 3
AS path: I
L-ISIS Preference: 14
Level: 2
Next hop type: Router, Next hop index: 0
Address: 0xb7a28f0
Next-hop reference count: 1
Next hop: 172.16.0.2 via ge-0/0/5.0, selected
Label operation: Push 800103
Label TTL action: prop-ttl
Load balance label: Label 800103: None;
Label element ptr: 0xb7a2880
Label parent element ptr: 0x0
Label element references: 1
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
State: Int
Inactive reason: Route Preference
Local AS:      11
Age: 45:40      Metric: 30

```



```

Validation State: unverified
ORR Generation-ID: 0
Task: IS-IS
AS path: I

127.0.0.2/8 (1 entry, 1 announced)
  *L-ISIS Preference: 14
    Level: 2
    Next hop type: Router, Next hop index: 0
    Address: 0xb7a29b0
    Next-hop reference count: 1
    Next hop: 172.16.0.2 via ge-0/0/5.0, selected
    Label operation: Push 800105
    Label TTL action: prop-ttl
    Load balance label: Label 800105: None;
    Label element ptr: 0xb7a2940
    Label parent element ptr: 0x0
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
    Session Id: 0x0
    State: Active Int
    Local AS: 11
    Age: 45:40 Metric: 20
    Validation State: unverified
    ORR Generation-ID: 0
    Task: IS-IS
    Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 3
    AS path: I

```

Meaning

Tunnel routes have been created for the PCE-controlled LSP destination with SPRING-TE as the protocol label.

Verifying Forwarding Table Entries

Purpose

Verify that the SPRING-TE LSP destination to Device R3 is installed in the forwarding table of the PCC.

Action

From operation mode, run the **show route forwarding-table destination ip-address extensive**

```
user@PCC> show route forwarding-table destination 127.0.0.3 extensive
```


Routing table: default.inet [Index 0]

Internet:

Enabled protocols: Bridging,

Destination: 127.0.0.3/8

Route type: user

Route reference: 0

Route interface-index: 0

Multicast RPF nh index: 0

P2mpidx: 0

Flags: sent to PFE, rt nh decoupled

Nexthop: 172.16.0.2

Next-hop type: unicast

Index: 581

Reference: 14

Next-hop interface: ge-0/0/5.0

Routing table: __pfe_private__.inet [Index 3]

Internet:

Enabled protocols: Bridging,

Destination: default

Route type: permanent

Route reference: 0

Route interface-index: 0

Multicast RPF nh index: 0

P2mpidx: 0

Flags: sent to PFE

Next-hop type: discard

Index: 517

Reference: 2

Routing table: __juniper_services__.inet [Index 5]

Internet:

Enabled protocols: Bridging,

Destination: default

Route type: permanent

Route reference: 0

Route interface-index: 0

Multicast RPF nh index: 0

P2mpidx: 0

Flags: sent to PFE

Next-hop type: discard

Index: 530

Reference: 2

Routing table: __master.anon__.inet [Index 6]

Internet:

Enabled protocols: Bridging, Dual VLAN,

Destination: default

Route type: permanent


```

Route reference: 0                      Route interface-index: 0
Multicast RPF nh index: 0
P2mpidx: 0
Flags: sent to PFE
Next-hop type: reject                  Index: 545      Reference: 1

```

Meaning

The SPRING-TE LSP destination IP address to Router R3 is installed as a forwarding entry.

Verifying Use of Tunnel Routes for Static Route Forwarding

Purpose

Verify that the static route is taking the tunnel route created for the SPRING-TE LSPs.

Action

From operational mode, run the **show route ip-address** and **show route forwarding-table destination ip-address** commands.

user@PCC>show route 100.1.1.1

```

inet.0: 17 destinations, 17 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100.1.1.1/32      *[Static/5] 00:33:36, metric2 0
                  > to 172.16.0.2 via ge-0/0/5.0, Push 16, Push 17(top)
                  to 172.16.0.2 via ge-0/0/5.0, Push 800103, Push 800105(top)

```

user@PCC> show route forwarding-table destination 100.1.1.1

```

Routing table: default.inet
Internet:
Enabled protocols: Bridging,
Destination      Type RtRef Next hop          Type Index   NhRef Netif
100.1.1.1/32     user    0          172.16.0.2      Push 16, Push 17(top) 590
2 ge-0/0/5.0

Routing table: __pfe_private__.inet
Internet:
Enabled protocols: Bridging,
Destination      Type RtRef Next hop          Type Index   NhRef Netif

```



```

default          perm      0          dscd      517      2

Routing table: __juniper_services__.inet
Internet:
Enabled protocols: Bridging,
Destination      Type RtRef Next hop          Type Index      NhRef Netif
default          perm      0          dscd      530      2

Routing table: __master.anon__.inet
Internet:
Enabled protocols: Bridging, Dual VLAN,
Destination      Type RtRef Next hop          Type Index      NhRef Netif
default          perm      0          rjct      545      1

```

Meaning

The outputs show that the static route to Router R3 uses the tunnel route created for the SPRING-TE LSP.

RELATED DOCUMENTATION

[Support of SPRING-TE for the Path Computation Element Protocol Overview | 1777](#)

Static Segment Routing Label Switched Path

IN THIS SECTION

- [Understanding Static Segment Routing LSP in MPLS Networks | 1812](#)
- [Example: Configuring Static Segment Routing Label Switched Path | 1836](#)

The segment routing architecture enables the ingress devices in a core network to steer traffic through explicit paths. You can configure these paths using segment lists to define the paths that the incoming traffic should take. The incoming traffic may be labeled or IP traffic, causing the forwarding operation at the ingress device to be either a label swap, or a destination-based lookup.

Understanding Static Segment Routing LSP in MPLS Networks

IN THIS SECTION

- [Introduction to Segment Routing LSPs | 1812](#)
- [Benefits of using Segment Routing LSPs | 1813](#)
- [Colored Static Segment Routing LSP | 1813](#)
- [Non-Colored Static Segment Routing LSP | 1814](#)
- [Static Segment Routing LSP Provisioning | 1821](#)
- [Static Segment Routing LSP Limitations | 1821](#)
- [Dynamic Creation of Segment Routing LSPs | 1822](#)
- [Color-Based Mapping of VPN Services | 1828](#)
- [Tunnel Templates for PCE-Initiated Segment Routing LSPs | 1835](#)

Source packet routing or segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network without relying on the intermediate nodes in the network to determine the actual path it should take.

Introduction to Segment Routing LSPs

Segment routing leverages the source routing paradigm. A device steers a packet through an ordered list of instructions, called segments. A segment can represent any instruction, topological or service-based. A segment can have a local semantic to a segment routing node or to a global node within a segment routing domain. Segment routing enforces a flow through any topological path and service chain while maintaining per-flow state only at the ingress device to the segment routing domain. Segment routing can be directly applied to the MPLS architecture with no change on the forwarding plane. A segment is encoded as an MPLS label. An ordered list of segments is encoded as a stack of labels. The segment to process is on the top of the stack. Upon completion of a segment, the related label is popped from the stack.

Segment routing LSPs can either be dynamic or static in nature.

Dynamic segment routing LSPs—When a segment routing LSP is created either by an external controller and downloaded to an ingress device through Path Computation Element Protocol (PCEP) extensions, or from a BGP segment routing policy through BGP segment routing extensions, the LSP is dynamically provisioned. The segment list of the dynamic segment routing LSP is contained in the PCEP Explicit Route Object (ERO), or the BGP segment routing policy of the LSP.

Static segment routing LSPs—When a segment routing LSP is created on the ingress device through local configuration, the LSP is statically provisioned.

A static segment routing LSP can further be classified as colored and non-colored LSPs based on the configuration of the **color** statement at the **[edit protocols source-packet-routing source-routing-path *lsp-name*]** hierarchy level.

For example:

```
[edit protocols]
  source-packet-routing {
    source-routing-path lsp_name {
      to destination_address;
      color color_value;
      binding-sid binding-label;
      primary segment_list_1_name weight weight;
      ...
      primary segment_list_n_name weight weight;
      secondary segment_list_n_name;
      sr-preference sr_preference_value;
    }
  }
```

Here, each primary and secondary statement refers to a segment list.

```
[edit protocols]
  source-packet-routing {
    segment-list segment_list_name {
      hop_1_name label sid_label;
      ...
      hop_n_name label sid_label;
    }
  }
```

Benefits of using Segment Routing LSPs

- Static segment routing does not rely on per LSP forwarding state on transit routers. Hence, removing the need of provisioning and maintaining per LSP forwarding state in the core.
- Provide higher scalability to MPLS networks.

Colored Static Segment Routing LSP

IN THIS SECTION

- [Understanding Colored Static Segment Routing LSPs | 1814](#)
- [Segment List of Colored Segment Routing LSPs | 1814](#)

A static segment routing LSP configured with the **color** statement is called a colored LSP.

Understanding Colored Static Segment Routing LSPs

Similar to a BGP segment routing policy, the ingress route of the colored LSP is installed in the **inetcolor.0** or **inet6color.0** routing tables, with **destination-ip-address, color** as key for mapping IP traffic.

A static colored segment routing LSP may have a binding SID, for which a route is installed in the **mpls.0** routing table. This binding SID label is used to map labeled traffic to the segment routing LSP. The gateways of the route are derived from the segment list configurations under the primary and secondary paths.

Segment List of Colored Segment Routing LSPs

The colored static segment routing LSPs already provide support for first hop label mode of resolving an LSP. However, first hop IP mode is not supported for colored segment routing LSPs. Starting in Junos OS Release 19.1R1, a commit check feature is introduced to ensure that all the segment lists contributing to the colored routes have the minimum label present for all hops. If this requirement is not met, the commit is blocked.

Non-Colored Static Segment Routing LSP

IN THIS SECTION

- [Understanding Non-Colored Segment Routing LSPs | 1814](#)
- [Segment List of Non-Colored Segment Routing LSPs | 1815](#)

A static segment routing LSP that is configured without the **color** statement is a non-colored LSP. Similar to PCEP segment routing tunnels, the ingress route is installed in the **inet.3** or **inet6.3** routing tables.

Junos OS supports non-colored static segment routing LSPs on ingress routers. You can provision non-colored static segment routing LSP by configuring one source routed path and one or more segment lists. These segment lists can be used by multiple non-colored segment routing LSPs.

Understanding Non-Colored Segment Routing LSPs

The non-colored segment routing LSP has a unique name and a destination IP address. An ingress route to the destination is installed in the **inet.3** routing table with a default preference of 8 and a metric of 1. This route allows non-colored services to be mapped to the segment routing LSP pertaining to the destination. In case the non-colored segment routing LSP does not require an ingress route then the ingress route can be disabled. A non-colored segment routing LSP uses binding SID label to achieve segment routing LSP stitching. This label that can be used to model the segment routing LSP as a segment that may be further used to construct other segment routing LSPs in a hierarchical manner. The transit of the binding SID label, by default, has a preference of 8 and a metric of 1.

Starting in Junos OS Release 18.2R1, statically configured non-colored segment routing LSPs on the ingress device are reported to the Path Computation Element (PCE) through a Path Computation Element Protocol

(PCEP) session. These non-colored segment routing LSPs may have binding service identifier (SID) labels associated with them. With this feature, the PCE can use this binding SID label in the label stack to provision PCE-initiated segment routing LSP paths.

A non-colored segment routing LSP can have a maximum of 8 primary paths. If there are multiple operational primary paths then the packet forwarding engine (PFE) distributes traffic over the paths based on the load balancing factors like the weight configured on the path. This is equal cost multi path (ECMP) if none of the paths have a weight configured on them or weighted ECMP if at least one of the paths has a non-zero weight configured on the paths. In both the cases, when one or some of the paths fail, the PFE rebalances the traffic over the remaining paths that automatically leads to achieving path protection. A non-colored segment routing LSP can have a secondary path for dedicated path protection. Upon failure of a primary path, the PFE rebalances the traffic to the remaining functional primary paths. Otherwise, the PFE switches the traffic to the backup path, hence achieving path protection. A non-colored segment routing LSP may specify a metric at `[edit protocols source-packet-routing source-routing-path lsp-name]` for its ingress and binding-SID routes. Multiple non-colored segment routing LSPs have the same destination address that contribute to the next hop of the ingress route.

Multiple non-colored segment routing LSPs have the same destination address that contribute to the next hop of the ingress route. Each path, either primary or secondary, of each segment routing LSP is considered as a gateway candidate, if the path is functional and the segment routing LSP has the best preference of all these segment routing LSPs. However, the maximum number of gateways that the next-hop can hold cannot exceed the RPD multi-path limit, which is 128 by default. Extra paths are pruned, firstly secondary paths and then primary paths. A given segment list may be referred multiple times as primary or secondary paths by these segment routing LSPs. In this case, there are multiple gateways, each having a unique segment routing LSP tunnel ID. These gateways are distinct, although they have identical outgoing label stack and interface. A non-colored segment routing LSP and a colored segment routing LSP may also have the same destination address. However, they correspond to different destination addresses for ingress routes, as the colored segment routing LSP's destination address is constructed with both its destination address and color.

NOTE: In the case where a static non-colored segment routing LSP and a PCEP-created segment routing LSP co-exist and have the same address that contributes to the same ingress route, if they also have the same preference. Otherwise, the segment routing LSP with the best preference is installed for the route.

Segment List of Non-Colored Segment Routing LSPs

A segment list consists of a list of hops. These hops are based on the SID label or an IP address. The number of SID labels in the segment list should not exceed the maximum segment list limit. You can configure the maximum segment list limit at the `[edit protocols source-packet-routing]` hierarchy level.

Prior to Junos OS Release 19.1R1, for a non-colored static segment routing LSP to be usable, the first hop of the segment list had to be an IP address of an outgoing interface and the second to *n*th hops could be

SID labels. Starting in Junos OS Release 19.1R1, this requirement does not apply, as the first hop of the non-colored static LSPs now provides support for SID labels, in addition to IP addresses. With the first hop label support, MPLS fast reroute (FRR) and weighted equal-cost multipath is enabled for resolving the static non-colored segment routing LSPs, similar to colored static LSPs.

For the first-hop label mode to take effect, you must include the **inherit-label-nexthops** statement globally or individually for a segment list, and the first hop of the segment list must include both IP address and label. If the first hop includes only IP address, the **inherit-label-nexthops** statement does not have any effect.

You can configure **inherit-label-nexthops** at any one of the following hierarchies. The **inherit-label-nexthops** statement takes effect only if the segment list first hop includes both IP address and label.

- **Segment list level**—At the `[edit protocols source-packet-routing segment-list segment-list-name]` hierarchy level.
- **Globally**—At the `[edit protocols source-packet-routing]` hierarchy level.

When the **inherit-label-nexthops** statement is configured globally, it takes precedence over the segment-list level configuration, and the **inherit-label-nexthops** configuration is applied to all the segment lists. When the **inherit-label-nexthops** statement is not configured globally, only segment lists with both labels and IP address present in the first hop, and configured with **inherit-label-nexthops** statement are resolved using SID labels.

For dynamic non-colored static LSPs, that is the PCEP-driven segment routing LSPs, the **inherit-label-nexthops** statement must be enabled globally, as the segment-level configuration is not applied.

[Table 16 on page 637](#) describes the mode of segment routing LSP resolution based on the first hop specification.

Table 49: Non-Colored Static LSP Resolution Based on First Hop Specification

First Hop Specification	Mode of LSP Resolution
IP address only For example: <pre> segment-list path-1 { hop-1 ip-address 172.0.12.2; hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	The segment list is resolved using the IP address.

Table 49: Non-Colored Static LSP Resolution Based on First Hop Specification (*continued*)

First Hop Specification	Mode of LSP Resolution
<p>SID only</p> <p>For example:</p> <pre>segment-list path-2 { hop-1 label 1000011; hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	<p>The segment list is resolved using SID labels.</p>
<p>IP address and SID (without the inherit-label-nexthops configuration)</p> <p>For example:</p> <pre>segment-list path-3 { hop1 { label 801006; ip-address 172.24.1.2; } hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	<p>By default, the segment list is resolved using IP address.</p>
<p>IP address and SID (with the inherit-label-nexthops configuration)</p> <p>For example:</p> <pre>segment-list path-3 { inherit-label-nexthops; hop1 { label 801006; ip-address 172.24.1.2; } hop-2 label 1000012; hop-3 label 1000013; hop-4 label 1000014; }</pre>	<p>The segment list is resolved using SID labels.</p>

You can use the **show route *ip-address* protocol spring-te active-path table inet.3** command to view the non-colored segment routing traffic-engineered LSPs having multiple segment lists installed in the inet.3 routing table.

For example:

user@host> **show route 7.7.7.7 protocol spring-te active-path table inet.3**

```
inet.3: 42 destinations, 59 routes (41 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

7.7.7.7/32      *[SPRING-TE/8] 00:01:25, metric 1, metric2 0
                > to 11.1.1.2 via et-0/0/0.1, Push 801007
                  to 21.1.1.2 via et-0/0/2.1, Push 801007
                  to 11.102.1.2 via et-0/0/0.2, Push 801007, Push 801002(top)

                  to 21.202.1.2 via et-0/0/2.2, Push 801007, Push 801005(top)

                  to 11.103.1.2 via et-0/0/0.3, Push 801007, Push 801003(top)

                  to 21.203.1.2 via et-0/0/2.3, Push 801007, Push 801006(top)

                  to 11.104.1.2 via et-0/0/0.4, Push 801007, Push 801003, Push
801002(top)
                  to 21.204.1.2 via et-0/0/2.4, Push 801007, Push 801006, Push
801005(top)
```


NOTE:

The first hop type of segment lists of a static segment routing LSP can cause a commit to fail, if:

- Different segment lists of a tunnel have different first hop resolution types. This is applicable to both colored and non-colored static segment routing LSPs. However, this does not apply for PCEP-driven LSPs; a system log message is generated for the mismatch in the first hop resolution type at the time of computing the path.

For example:

```
segment-list path-1 {
  hop-1 ip-address 172.0.12.2;
  hop-2 label 1000012;
  hop-3 label 1000013;
  hop-4 label 1000014;
}
segment-list path-2 {
  hop-1 label 1000011;
  hop-2 label 1000012;
  hop-3 label 1000013;
  hop-4 label 1000014;
}
source-routing-path lsp1 {
  to 172.10.10.1;
  primary {
    path-1;
    path-2;
  }
}
```

The commit of tunnel *lsp1* fails, as path-1 is of IP address mode and path-2 is of label mode.

- The binding SID is enabled for the static non-colored LSP whose segment list type is SID label.

For example:

```
segment-list path-3 {
  hop-1 label 1000011;
  hop-2 label 1000012;
  hop-3 label 1000013;
  hop-4 label 1000014;
}
source-routing-path lsp1 {
```



```

to 172.10.10.1;
binding-sid 333;
primary {
    path-3;
}
}

```

Configuring binding SID over label segment list is supported only for colored static LSPs and not for no-colored static LSPs.

Static Segment Routing LSP Provisioning

Segment provisioning is performed on per-router basis. For a given segment on a router, a unique service identifier (SID) label is allocated from a desired label pool which may be from the dynamic label pool for an adjacency SID label or from the segment routing global block (SRGB) for a prefix SID or node SID. The adjacency SID label can be dynamically allocated, which is the default behavior, or be allocated from a local static label pool (SRLB). A route for the SID label is then installed in the mpls.0 table.

Junos OS allows static segment routing LSPs by configuring the **segment** statement at the **[edit protocols mpls static-label-switched-path static-label-switched-path]** hierarchy level. A static segment LSP is identified by a unique SID label that falls under Junos OS static label pool. You can configure the Junos OS static label pool by configuring the **static-label-range static-label-range** statement at the **[edit protocols mpls label-range]** hierarchy level.

Static Segment Routing LSP Limitations

- Junos OS currently has a limitation that the next hop cannot be built to push more than the maximum segment list depth labels. So, a segment list with more than the maximum SID labels (excluding the SID label of the first hop which is used to resolve forwarding next-hop) is not usable for colored or non-colored segment routing LSPs. Also, the actual number allowed for a given segment routing LSP may be even lower than the maximum limit, if an MPLS service is on the segment routing LSP or the segment routing LSP is on a link or a node protection path. In all cases, the total number of service labels, SID labels, and link or node protection labels must not exceed the maximum segment list depth. You can configure the maximum segment list limit at **[edit protocols source-packet-routing]** hierarchy level. Multiple non-colored segment routing LSPs with less than or equal to the maximum SID labels can be stitched together to construct a longer segment routing LSP. This is called segment routing LSP stitching. It can be achieved using binding-SID label.
- The segment routing LSP stitching is actually performed at path level. If a non-colored segment routing LSP has multiple paths that is multiple segment lists, each path can be independently stitched to another non-colored segment routing LSP at a stitching point. A non-colored segment routing LSP which is dedicated to stitching may disable ingress route installation by configuring **no-ingress** statement at **[edit protocols source-packet-routing source-routing-path lsp-name]** hierarchy level.

- A maximum of 8 primary paths and 1 secondary path are supported per non-colored static segment routing LSP. If there is a violation in configuration, commit check fails with an error.
- If any segment-list is configured with more labels than the maximum segment list depth, the configuration commit check fails with an error.

Dynamic Creation of Segment Routing LSPs

IN THIS SECTION

- [Configuring Dynamic Segment Routing LSP Template | 1822](#)
- [Resolving Dynamic Segment Routing LSPs | 1823](#)
- [Considerations for Configuring Dynamic Creation of Segment Routing LSPs | 1827](#)
- [Services Supported over Dynamic Segment Routing LSPs | 1827](#)
- [Behavior With Multiple Tunnel Sources in Segment Routing | 1828](#)
- [Dynamic Segment Routing LSPs Limitations | 1828](#)

In segment routing networks that have each provider edge (PE) device connected to every other PE device, a large amount of configuration is required for setting up the segment routing label-switched paths (LSPs), although only a few segment routing traffic-engineered (SR-TE) paths may be used. You can enable BGP-triggered dynamic creation of these LSPs to reduce the amount of configuration in such deployments.

Configuring Dynamic Segment Routing LSP Template

To configure the template for enabling dynamic creation of segment routing LSPs, you must include the [spring-te](#) statement at the `[edit routing-options dynamic-tunnels]` hierarchy.

- The following is a sample configuration for color dynamic segment routing LSP template:

```
[edit routing-options]
dynamic-tunnels {
  <dynamic-tunnel-name> {
    spring-te {
      source-routing-path-template {
        <template-name1> color [c1 c2];
        <template-name2> color [c3];
        <template-name3> color-any;
      }
      destination-networks {
        <dest1>;
        <dest2>;
      }
    }
  }
}
```



```

    }
  }
}

```

- The following is a sample configuration for non-color dynamic segment routing LSP template:

```

dynamic-tunnels {
  <dynamic-tunnel-name> {
    spring-te {
      source-routing-path-template {
        <template-name1>;
      }
      destination-networks {
        <dest1>;
        <dest2>;
      }
    }
  }
}

```

Resolving Dynamic Segment Routing LSPs

IN THIS SECTION

- [Resolving Colored Dynamic Segment Routing LSP | 1823](#)
- [Resolving Uncolored Dynamic Segment Routing LSPs | 1824](#)
- [Dynamic Segment Routing LSP Sample Configuration | 1824](#)

Resolving Colored Dynamic Segment Routing LSP

When the BGP prefixes are assigned with color community, they initially get resolved over the catch-all-route-for-that-particular-color policy, and in turn, the SR-TE template on which the BGP prefix should be resolved onto is identified. The destinations SID is then derived from the BGP payload prefix next-hop attribute. For example, if the next hop of the BGP payload prefix is an IP address that belongs to Device A, then the node-SID of Device A is taken and a corresponding label is prepared and pushed to the bottom of the stack. The BGP payload prefix is resolved in a color-only mode, where the node-SID of Device A is at the bottom of the final label stack, and the SR-TE path labels are on top.

The final LSP template name is a combination of prefix, color, and tunnel name; for example, **<prefix>:<color>:dt-srte-<tunnel-name>**. The color in the LSP name is displayed in hexadecimal format, and the format of the tunnel name is similar to that of RSVP-triggered tunnel LSP names.

To successfully resolve a colored destination network, the color should have a valid template mapping, either to a specific color, or through the **color-any** template. Without a valid mapping, the tunnel is not created and the BGP route requesting for resolution remains unresolved.

Resolving Uncolored Dynamic Segment Routing LSPs

The catch-all routes for non-colored LSPs are added to the inet.3 routing table. The non-colored tunnel destination must be configured in a different **spring-te** configuration with only one template name in the mapping list. This template name is used for all the tunnel routes matching any of the destination networks configured under the same **spring-te** configuration. These tunnels are similar to RSVP tunnels in functionality.

The final LSP template name is a combination of prefix and tunnel name; for example, **<prefix>:dt-srte-<tunnel-name>**.

Dynamic Segment Routing LSP Sample Configuration

IN THIS SECTION

- [Colored Dynamic Segment Routing LSPs | 1824](#)
- [Non-Colored Dynamic Segment Routing LSPs | 1825](#)
- [Unresolved Dynamic Segment Routing LSP | 1826](#)

The dynamic segment routing LSP template always carries a partial path. The last hop node SID is derived automatically at the tunnel creation time depending on the protocol next-hop address (PNH) node SID. The same template can be used by multiple tunnels to different destinations. In such cases, the partial path remains the same, and only the last hop changes depending on the PNH. Dynamic segment routing LSP templates are not common to a single tunnel, as a result a full path cannot be carried on it. You can use a segment list if a full path is to be used.

Colored Dynamic Segment Routing LSPs

Sample configuration for colored dynamic segment routing LSPs:

```
protocols source-packet-routing {
  source-routing-path-template sr_lsp1 {
    primary sr_sl1
    primary sr_sl2 weight 2
    sr-preference 180;
  }
}
dynamic-tunnels tunnel1 {
  spring-te {
    source-routing-path-template {
```



```

        sr_lsp1 color [ 123 124 125 ];
        sr_lsp2 color-any
    }
    destination-networks {
        22.33.44.0/24;
    }
}
}
}

```

For the above-mentioned sample configuration, the route entries are as follows:

inetcolor.0 tunnel route: 22.33.44.0-0/24 --> RT_NH_TUNNEL

inetcolor6.0 tunnel route: ffff::22.33.44.0-0/120 --> RT_NH_TUNNEL

BGP prefix to tunnel mapping:

R1(prefix) -> 22.33.44.55-101(PNH) LSP tunnel name = 22.33.44.55:65:dt-srte-tunnel1

R1(prefix) -> ffff::22.33.44.55-101(PNH) LSP tunnel name = 22.33.44.55:65:dt-srte-tunnel1

R1(prefix) -> ffff::22.33.44.55-124(PNH) LSP tunnel name = 22.33.44.55:7c:dt-srte-tunnel1

inetcolor.0 tunnel route:

22.33.44.55-101/64 --> <next-hop>

22.33.44.55-124/64 --> <next-hop>

inetcolor6.0 tunnel route:

ffff::22.33.44.55-101/160 --> <next-hop>

ffff::22.33.44.55-124/160 --> <next-hop>

Non-Colored Dynamic Segment Routing LSPs

Sample configuration for non-colored dynamic segment routing LSPs:

```

protocols source-packet-routing {
    source-routing-path-template sr_lsp1 {
        primary sr_sl1
        primary sr_sl2 weight 2
        sr-preference 180;
    }
}
dynamic-tunnels {
    tunnel1 {
        spring-te {
            source-routing-path-template {

```



```

        sr_lsp1 color [ 101 ];
    }
    destination-networks {
        22.33.44.0/24;
    }
}
}
tunnel2 {
    spring-te {
        source-routing-path-template {
            sr_lsp1;
        }
        destination-networks {
            22.33.44.0/24;
        }
    }
}
}
}

```

For the above-mentioned sample configuration, the route entries are as follows:

inet.3 tunnel route: 22.33.44.0/24 --> RT_NH_TUNNEL

inet6.3 tunnel route: ffff::22.33.44.0/120 --> RT_NH_TUNNEL

BGP prefix to tunnel mapping:

R1(prefix) -> 22.33.44.55(PNH) LSP template name = LSP1 --- 22.33.44.55:dt-srte-tunnel2

R1(prefix) -> ffff::22.33.44.55(PNH) LSP template name = LSP1 --- 22.33.44.55:dt-srte-tunnel2

inet.3 tunnel route: 22.33.44.55/32 --> <next-hop>

inet6.3 tunnel route: ffff::22.33.44.55/128 --> <next-hop>

Unresolved Dynamic Segment Routing LSP

Sample configuration for unresolved dynamic segment routing LSPs:

```

protocols source-packet-routing {
    source-routing-path-template sr_lsp1 {
        primary sr_sl1
        primary sr_sl2 weight 2
        sr-preference 180;
    }
}
dynamic-tunnels tunnel1 {
    spring-te {

```



```

source-routing-path-template {
    sr_lsp1 color [120 121 122 123];
}
destination-networks {
    22.33.44.0/24;
    1.1.1.0/24;
}
}
}

```

For the above-mentioned sample configuration, the route entries are as follows:

inetcolor.0 tunnel route: 22.33.44.0 - 0/24 --> RT_NH_TUNNEL 1.1.1.0 - 0 /24 --> RT_NH_TUNNEL

inetcolor6.0 tunnel route: ffff::22.33.44.0 - 0/120 --> RT_NH_TUNNEL ffff::1.1.1.0 - 0 /24 --> RT_NH_TUNNEL

BGP prefix to tunnel mapping: R1(prefix) -> 22.33.44.55-124(PNH) Tunnel will not be created. (Template not found for the color).

Considerations for Configuring Dynamic Creation of Segment Routing LSPs

When configuring the dynamic creation of segment routing LSPs, take the following into consideration:

- A template can be assigned with a color object. When the dynamic tunnel **spring-te** configuration includes a template with a color object, you must configure all other templates with color objects as well. All destinations are assumed to be colored within that configuration.
- A template can have a list of colors defined on it, or can be configured with the **color-any** option. Both these options can coexist in the same **spring-te** configuration. In such cases, templates assigned with specific colors have a higher preference.
- In a **spring-te** configuration, only one template can be defined with the **color-any** option.
- The color-to-template mapping is done on a one-to-one basis. One color cannot map to multiple templates.
- The template name should be configured in the **spring-te** statement under the **[edit protocols]** hierarchy, and should have the **primary** option enabled.
- Colored and non-colored destinations cannot co-exist in the same **spring-te** configuration.
- You cannot configure same destination networks, with or without color, under different **spring-te** configuration statements.
- In non-colored **spring-te** configuration, only one template can be configured without color object.

Services Supported over Dynamic Segment Routing LSPs

The following services are supported over colored dynamic segment routing LSPs:

- Layer 3 VPN

- BGP EVPN
- Export policy services

The following services are supported over non-colored dynamic segment routing LSPs:

- Layer 3 VPN
- Layer 2 VPN
- Multipath configurations

Behavior With Multiple Tunnel Sources in Segment Routing

When two sources download routes to the same destination from segment routing (for example static and dynamic sourced tunnels), then the segment routing preference is used for choosing the active route entry. A higher value has greater preference. In case the preference remains the same, then the tunnel source is used to determine the route entry.

Dynamic Segment Routing LSPs Limitations

The dynamic SR-TE LSPs do not support the following features and functionalities:

- IPv6 segment routing tunnels.
- Static tunnels.
- 6PE is not supported.
- Distributed CSPF.
- sBFD and LDP tunnelling is not supported for dynamic SR-TE LSPs and in a template.
- Install and B-SID routes in a template.

Color-Based Mapping of VPN Services

IN THIS SECTION

- [VPN Service Coloring | 1829](#)
- [Specifying VPN Service Mapping Mode | 1832](#)
- [Color-IP Protocol Next Hop Resolution | 1833](#)
- [Fallback to IP Protocol Next Hop Resolution | 1834](#)
- [Supported and Unsupported Features for Color-Based Mapping of VPN Services | 1834](#)

You can specify color as a protocol next hop constraint (in addition to the IPv4 or IPv6 address) for resolving transport tunnels over static colored and BGP segment routing traffic-engineered (SRTE) LSPs. This is called the color-IP protocol next hop resolution, where you are required to configure a resolution-map

and apply to the VPN services. With this feature, you can enable color-based traffic steering of Layer 2 and Layer 3 VPN services.

Junos OS supports colored SRTE LSPs associated with a single color. The color-based mapping of VPN services feature is supported on static colored LSPs and BGP SRTE LSPs.

VPN Service Coloring

IN THIS SECTION

- [Egress Color Assignment | 1829](#)
- [Ingress Color Assignment | 1831](#)

In general, a VPN service may be assigned a color on the egress router where the VPN NLRI is advertised, or on an ingress router where the VPN NLRI is received and processed.

You can assign a color to the VPN services at different levels:

- Per routing instance.
- Per BGP group.
- Per BGP neighbor.
- Per prefix.

Once you assign a color, the color is attached to a VPN service in the form of BGP color extended community.

You can assign multiple colors to a VPN service, referred to as multi-color VPN services. In such cases, the last color attached is considered as the color of the VPN service, and all other colors are ignored.

Multiple colors are assigned by egress and/or ingress devices through multiple policies in the following order:

- BGP export policy on the egress device.
- BGP import policy on the ingress device.
- VRF import policy on the ingress device.

The two modes of VPN service coloring are:

Egress Color Assignment

In this mode, the egress device (that is, the advertiser of the VPN NLRI) is responsible for coloring the VPN service. To enable this mode, you can define a routing policy, and apply it in the VPN service's

routing-instance **vrf-export**, group export, or group neighbor export at the **[edit protocols bgp]** hierarchy level. The VPN NLRI is advertised by BGP with the specified color extended community.

For example:

```
[edit routing-options]
community red-comm {
  members color:0:50;
}
```

```
[edit policy-options]
policy-statement pol-color {
  term t1 {
    from {
      [any match conditions];
    }
    then {
      community add red-comm;
      accept;
    }
  }
}
```

```
[edit routing-instances]
vpn-X {
  ...
  vrf-export pol-color ...;
}
```

Or

NOTE: When you apply the routing policy as an export policy of a BGP group or BGP neighbor, you must include the **vpn-apply-export** statement at the BGP, BGP group, or BGP neighbor level in order for the policy to take an effect on the VPN NLRI.

```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-A {
    export pol-color ...;
```



```

        vpn-apply-export;
    }
}

```

The routing policies are applied to Layer 3 VPN prefix NLRIs, Layer 2 VPN NLRIs, and EVPN NLRIs. The color extended community is inherited by all the VPN routes, imported, and installed in the target VRFs on one or multiple ingress devices.

Ingress Color Assignment

In this mode, the ingress device (that is, the receiver of the VPN NLRI) is responsible for coloring the VPN service. To enable this mode, you can define a routing policy, and apply it to the VPN service's routing-instance **vrf-import**, group import, or group neighbor import at the **[edit protocols bgp]** hierarchy level. All the VPN routes matching the routing policy is attached with the specified color extended community.

For example:

```

[edit routing-options]
community red-comm {
    members color:0:50;
}

```

```

[edit policy-options]
policy-statement pol-color {
    term t1 {
        from {
            [any match conditions];
        }
        then {
            community add red-comm;
            accept;
        }
    }
}

```

```

[edit routing-instances]
vpn-Y {
    ...
    vrf-import pol-color ...;
}

```

Or


```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-B {
    import pol-color ...;
  }
}
```

Specifying VPN Service Mapping Mode

To specify flexible VPN service mapping modes, you must define a policy using the **resolution-map** statement, and refer the policy in a VPN service's routing-instance **vrf-import**, group import, or group neighbor import at the **[edit protocols bgp]** hierarchy level. All the VPN routes matching the routing policy are attached with the specified resolution-map.

For example:

```
[edit policy-options]
resolution-map map-A {
  <mode-1>;
  <mode-2>;
  ...
}
policy-statement pol-resolution {
  term t1 {
    from {
      [any match conditions];
    }
    then {
      resolution-map map-A;
      accept;
    }
  }
}
```

You can apply import policy to the VPN service's routing-instance.

```
[edit routing-instances]
vpn-Y {
  ...
  vrf-import pol-resolution ...;
}
```

You can also apply the import policy to a BGP group or BGP neighbor.


```
[edit protocols bgp]
group PEs {
  ...
  neighbor PE-B {
    import pol-resolution ...;
  }
}
```

NOTE: Each VPN service mapping mode should have a unique name defined in the resolution-map. Only a single entry of IP-color is supported in the resolution-map, where the VPN route(s) are resolved using a colored-IP protocol next hop in the form of **ip-address:color**.

Color-IP Protocol Next Hop Resolution

The protocol next hop resolution process is enhanced to support colored-IP protocol next hop resolution. For a colored VPN service, the protocol next hop resolution process takes a color and a resolution-map, builds a colored-IP protocol next hop in the form of *IP-address:color*, and resolves the protocol next hop in the inet6color.0 routing table.

You must configure a policy to support multipath resolution of colored Layer 2 VPN, Layer 3 VPN, or EVPN services over colored LSPs. The policy must then be applied with the relevant RIB table as the resolver import policy.

For example:

```
[edit policy-options]
policy-statement mpath {
  then multipath-resolve;
}
```

```
[edit routing-options]
resolution {
  rib bgp.l3vpn.0 {
    inetcolor-import mpath;
  }
}
resolution {
  rib bgp.l3vpn-inet6.0 {
    inet6color-import mpath;
  }
}
```



```

resolution {
    rib bgp.l2vpn.0 {
        inetcolor-import mpath;
    }
}
resolution {
    rib mpls.0 {
        inetcolor-import mpath;
    }
}
resolution {
    rib bgp.evpn.0 {
        inetcolor-import mpath;
    }
}

```

Fallback to IP Protocol Next Hop Resolution

If a colored VPN service does not have a resolution-map applied to it, the VPN service ignores its color and falls back to the IP protocol next hop resolution. Conversely, if a non-colored VPN service has a resolution-map applied to it, the resolution-map is ignored, and the VPN service uses the IP protocol next hop resolution.

The fallback is a simple process from colored SRTE LSPs to LDP LSPs, by using a RIB group for LDP to install routes in inet{6}color.0 routing tables. A longest prefix match for a colored-IP protocol next hop ensures that if a colored SRTE LSP route does not exist, an LDP route with a matching IP address should be returned.

Supported and Unsupported Features for Color-Based Mapping of VPN Services

The following features and functionality are supported with color-based mapping of VPN services:

- BGP Layer 3 VPN
- BGP Layer 2 VPN (Kompella Layer 2 VPN)
- BGP EVPN
- Resolution-map with a single IP-color option.
- Colored IPv4 and IPv6 protocol next hop resolution.
- Routing information base (also known as routing table) group based fallback to LDP LSP in inetcolor.0 routing table.
- Colored SRTE LSP.
- Virtual platforms.

- 64-bit Junos OS.
- Logical systems.

The following features and functionality are not supported with color-based mapping of VPN services:

- Colored MPLS LSPs, such as RSVP, LDP, BGP-LU, static.
- Layer 2 circuit
- FEC-129 BGP auto-discovered and LDP-signaled Layer 2 VPN.
- VPLS
- MVPN
- IPv4 and IPv6 using resolution-map.
- BGP labeled unicast.

Tunnel Templates for PCE-Initiated Segment Routing LSPs

You can configure a tunnel template for PCE-initiated segment routing LSPs to pass down two additional parameters for these LSPs - Bidirectional forwarding detection (BFD) and LDP tunneling.

When a PCE-Initiated segment routing LSP is being created, the LSP is checked against policy statements (if any) and if there is a match, the policy applies the configured template for that LSP. The template configuration is inherited only if it is not provided by the LSP source (PCEP); for example, metric.

To configure a template:

1. Include the [source-routing-path-template](#) statement at the **[edit protocols source-packet-routing]** hierarchy level. You can configure the additional BFD and LDP tunneling parameters here.
2. Include the [source-routing-path-template-map](#) statement at the **[edit protocols source-packet-routing]** hierarchy level to list the policy statements against which the PCE-initiated LSP should be checked.
3. Define a policy to list the LSPs on which the template should be applied.

The **from** statement can include either the LSP name or LSP regular expression using the **lsp** and **lsp-regex** match conditions. These options are mutually exclusive, so you can specify only one option at a given point in time.

The **then** statement must include the **sr-te-template** option with an accept action. This applies the template to the PCE-initiated LSP.

Take the following into consideration when configuring a template for PCE-initiated LSPs:

- Template configuration is not applicable to statically configured segment routing LSPs, or any other client's segment routing LSP.
- PCEP-provided configuration has precedence over template configuration.

- PCEP LSP does not inherit template segment-list configuration.

Example: Configuring Static Segment Routing Label Switched Path

IN THIS SECTION

- [Requirements | 1836](#)
- [Overview | 1836](#)
- [Configuration | 1837](#)
- [Verification | 1850](#)

This example shows how to configure static segment routing label switched paths (LSPs) in MPLS networks. This configuration helps to bring higher scalability to MPLS networks.

Requirements

This example uses the following hardware and software components:

- Seven MX Series 5G Universal Routing Platforms
- Junos OS Release 18.1 or later running on all the routers

Before you begin, be sure you configure the device interfaces.


Overview

Junos OS a set of explicit segment routing paths are configured on the ingress router of a non-colored static segment routing tunnel by configuring the **segment-list** statement at the **[edit protocols source-packet-routing]** hierarchy level. You can configure segment routing tunnel by configuring the **source-routing-path** statement at **[edit protocols source-packet-routing]** hierarchy level. The segment routing tunnel has a destination address and one or more primary paths and optionally secondary paths that refer to the segment list. Each segment list consists of a sequence of hops. For non-colored static segment routing tunnel, the first hop of the segment list specifies an immediate next hop IP address and the second to Nth hop specifies the segment identifies (SID) labels corresponding to the link or node which the path traverses. The route to the destination of the segment routing tunnel is installed in inet.3 table.

Topology

In this example, configure layer 3 VPN on the provider edge routers PE1 and PE5. Configure the MPLS protocol on all the routers. The segment routing tunnel is configured from router PE1 to router PE5 with a primary path configured on router PE1 and router PE5 . Router PE1 is also configured with secondary path for path protection. The transit routers PE2 to PE4 are configured with adjacency SID labels with label pop and an outgoing interface.

Figure 144: Static Segment Routing Label Switched Path



Configuration

IN THIS SECTION

- Configuring Device PE1 | 1841
- Configuring Device PE2 | 1847
- Results | 1848

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

PE1

```
set interfaces ge-0/0/0 unit 0 family inet address 10.10.12.1/24
set interfaces ge-0/0/0 unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/1 unit 0 family inet address 10.10.13.1/24
set interfaces ge-0/0/1 unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/5 unit 0 family inet address 10.10.17.1/24
set routing-options autonomous-system 65000
set routing-options forwarding-table export load-balance-policy
set routing-options forwarding-table chained-composite-next-hop ingress l3vpn
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols bgp group pe type internal
set protocols bgp group pe local-address 192.168.147.211
set protocols bgp group pe family inet-vpn unicast
set protocols bgp group pe neighbor 192.168.146.181
```



```

set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols source-packet-routing segment-list sl-15-primary hop-1 ip-address 10.10.13.3
set protocols source-packet-routing segment-list sl-15-primary hop-2 label 1000134
set protocols source-packet-routing segment-list sl-15-primary hop-3 label 1000145
set protocols source-packet-routing segment-list sl-15-backup hop-1 ip-address 10.10.12.2
set protocols source-packet-routing segment-list sl-15-backup hop-2 label 1000123
set protocols source-packet-routing segment-list sl-15-backup hop-3 label 1000134
set protocols source-packet-routing segment-list sl-15-backup hop-4 label 1000145
set protocols source-packet-routing source-routing-path lsp-15 to 192.168.146.181
set protocols source-packet-routing source-routing-path lsp-15 binding-sid 1000999
set protocols source-packet-routing source-routing-path lsp-15 primary sl-15-primary
set protocols source-packet-routing source-routing-path lsp-15 secondary sl-15-backup
set policy-options policy-statement VPN-A-export term a from protocol ospf
set policy-options policy-statement VPN-A-export term a from protocol direct
set policy-options policy-statement VPN-A-export term a then community add VPN-A
set policy-options policy-statement VPN-A-export term a then accept
set policy-options policy-statement VPN-A-export term b then reject
set policy-options policy-statement VPN-A-import term a from protocol bgp
set policy-options policy-statement VPN-A-import term a from community VPN-A
set policy-options policy-statement VPN-A-import term a then accept
set policy-options policy-statement VPN-A-import term b then reject
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf from route-filter 10.10.0.0/16 orlonger
set policy-options policy-statement bgp-to-ospf then accept
set policy-options policy-statement load-balance-policy then load-balance per-packet
set policy-options community VPN-A members target:65000:1
set routing-instances VRF1 instance-type vrf
set routing-instances VRF1 interface ge-0/0/5.0
set routing-instances VRF1 route-distinguisher 192.168.147.211:1
set routing-instances VRF1 vrf-import VPN-A-import
set routing-instances VRF1 vrf-export VPN-A-export
set routing-instances VRF1 vrf-table-label
set routing-instances VRF1 protocols ospf export bgp-to-ospf
set routing-instances VRF1 protocols ospf area 0.0.0.0 interface ge-0/0/5.0

```

PE2

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.12.2/24

```



```

set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.23.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set protocols mpls static-label-switched-path adj-23 segment 1000123
set protocols mpls static-label-switched-path adj-23 segment next-hop 10.10.23.3
set protocols mpls static-label-switched-path adj-23 segment pop
set protocols mpls static-label-switched-path adj-21 segment 1000221
set protocols mpls static-label-switched-path adj-21 segment next-hop 10.10.12.1
set protocols mpls static-label-switched-path adj-21 segment pop
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0

```

PE3

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.13.3/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.23.3/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.10.34.3/24
set interfaces ge-0/0/2 unit 0 family mpls
set protocols mpls static-label-switched-path adj-34 segment 1000134
set protocols mpls static-label-switched-path adj-34 segment next-hop 10.10.34.4
set protocols mpls static-label-switched-path adj-34 segment pop
set protocols mpls static-label-switched-path adj-32 segment 1000232
set protocols mpls static-label-switched-path adj-32 segment next-hop 10.10.23.2
set protocols mpls static-label-switched-path adj-32 segment pop
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0

```

PE4


```

set interfaces ge-0/0/2 unit 0 family inet address 10.10.34.4/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 10.10.45.4/24
set interfaces ge-0/0/3 unit 0 family mpls
set protocols mpls static-label-switched-path adj-45 segment 1000145
set protocols mpls static-label-switched-path adj-45 segment next-hop 10.10.45.5
set protocols mpls static-label-switched-path adj-45 segment pop
set protocols mpls static-label-switched-path adj-43 segment 1000243
set protocols mpls static-label-switched-path adj-43 segment next-hop 10.10.34.3
set protocols mpls static-label-switched-path adj-43 segment pop
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0

```

PE5

```

set interfaces ge-0/0/3 unit 0 family inet address 10.10.45.5/24
set interfaces ge-0/0/3 unit 0 family mpls maximum-labels 5
set interfaces ge-0/0/4 unit 0 family inet address 10.10.56.5/24
set routing-options autonomous-system 65000
set protocols mpls interface ge-0/0/3.0
set protocols mpls label-range static-label-range 1000000 1000999
set protocols bgp group pe type internal
set protocols bgp group pe local-address 192.168.146.181
set protocols bgp group pe family inet-vpn unicast
set protocols bgp group pe neighbor 192.168.147.211
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols bfd sbfd local-discriminator 0.0.0.32 minimum-receive-interval 1000
set protocols source-packet-routing segment-list sl-51 hop-1 ip-address 10.10.45.4
set protocols source-packet-routing segment-list sl-51 hop-2 label 1000243
set protocols source-packet-routing segment-list sl-51 hop-3 label 1000232
set protocols source-packet-routing segment-list sl-51 hop-4 label 1000221
set protocols source-packet-routing source-routing-path lsp-51 to 192.168.147.211
set protocols source-packet-routing source-routing-path lsp-51 primary sl-51
set policy-options policy-statement VPN-A-export term a from protocol ospf
set policy-options policy-statement VPN-A-export term a from protocol direct
set policy-options policy-statement VPN-A-export term a then community add VPN-A

```



```

set policy-options policy-statement VPN-A-export term a then accept
set policy-options policy-statement VPN-A-export term b then reject
set policy-options policy-statement VPN-A-import term a from protocol bgp
set policy-options policy-statement VPN-A-import term a from community VPN-A
set policy-options policy-statement VPN-A-import term a then accept
set policy-options policy-statement VPN-A-import term b then reject
set policy-options policy-statement bgp-to-ospf from protocol bgp
set policy-options policy-statement bgp-to-ospf from route-filter 10.10.0.0/16 orlonger
set policy-options policy-statement bgp-to-ospf then accept
set policy-options community VPN-A members target:65000:1
set routing-instances VRF1 instance-type vrf
set routing-instances VRF1 interface ge-0/0/4.0
set routing-instances VRF1 route-distinguisher 192.168.146.181:1
set routing-instances VRF1 vrf-import VPN-A-import
set routing-instances VRF1 vrf-export VPN-A-export
set routing-instances VRF1 vrf-table-label
set routing-instances VRF1 protocols ospf export bgp-to-ospf
set routing-instances VRF1 protocols ospf area 0.0.0.0 interface ge-0/0/4.0

```

CE1

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.17.7/24
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0

```

CE2

```

set interfaces ge-0/0/4 unit 0 family inet address 10.10.56.6/24
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0

```

Configuring Device PE1

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE1:

1. Configure the interfaces.


```
[edit interfaces]
set ge-0/0/0 unit 0 family inet address 10.10.12.1/24
set ge-0/0/0 unit 0 family mpls maximum-labels 5

set ge-0/0/1 unit 0 family inet address 10.10.13.1/24
set ge-0/0/1 unit 0 family mpls maximum-labels 5

set ge-0/0/5 unit 0 family inet address 10.10.17.1/24
```

2. Configure autonomous system number and options to control packet forwarding routing options.

```
[edit routing-options]
set autonomous-system 65000
set forwarding-table export load-balance-policy
set forwarding-table chained-composite-next-hop ingress l3vpn
```

3. Configure the interfaces with the MPLS protocol and configure the MPLS label range.

```
[edit protocols mpls]
set interface ge-0/0/0.0
set interface ge-0/0/1.0
set label-range static-label-range 1000000 1000999
```

4. Configure the type of peer group, local address, protocol family for NLRI in updates, and IP address of a neighbor for the peer group.

```
[edit protocols bgp]
set group pe type internal
set group pe local-address 192.168.147.211
set group pe family inet-vpn unicast
set group pe neighbor 192.168.146.181
```

5. Configure the protocol area interfaces.

```
[edit protocols ospf]
set area 0.0.0.0 interface ge-0/0/0.0
set area 0.0.0.0 interface lo0.0
set area 0.0.0.0 interface ge-0/0/1.0
```


6. Configure the IPv4 address and labels of primary and secondary paths for source routing-traffic engineering (TE) policies of protocol source packet routing (SPRING).

```
[edit protocols source-packet-routing segment-list]
set sl-15-primary hop-1 ip-address 10.10.13.3
set sl-15-primary hop-2 label 1000134
set sl-15-primary hop-3 label 1000145
set sl-15-backup hop-1 ip-address 10.10.12.2
set sl-15-backup hop-2 label 1000123
set sl-15-backup hop-3 label 1000134
set sl-15-backup hop-4 label 1000145
```

7. Configure destination IPv4 address, binding SID label, primary, and secondary source routing path for protocol SPRING.

```
[edit protocols source-packet-routing source-routing-path]
set lsp-15 to 192.168.146.181
set lsp-15 binding-sid 1000999
set lsp-15 primary sl-15-primary
set lsp-15 secondary sl-15-backup
```

8. Configure policy options.

```
[edit policy-options policy-statement]
set VPN-A-export term a from protocol ospf
set VPN-A-export term a from protocol direct
set VPN-A-export term a then community add VPN-A
set VPN-A-export term a then accept
set VPN-A-export term b then reject
set VPN-A-import term a from protocol bgp
set VPN-A-import term a from community VPN-A
set VPN-A-import term a then accept
set VPN-A-import term b then reject
set bgp-to-ospf from protocol bgp
set bgp-to-ospf from route-filter 10.10.0.0/16 orlonger
set bgp-to-ospf then accept
set load-balance-policy then load-balance per-packet
```

9. Configure BGP community information.

```
[edit policy-options]
```



```
set community VPN-A members target:65000:1
```

10. Configure routing instance VRF1 with instance type, interface, router distinguisher, VRF import, export and table label. Configure export policy and interface of area for protocol OSPF.

```
[edit routing-instances VRF1]
set instance-type vrf
set interface ge-0/0/5.0
set route-distinguisher 192.168.147.211:1
set vrf-import VPN-A-import
set vrf-export VPN-A-export
set vrf-table-label
set protocols ospf export bgp-to-ospf
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, **show routing-options**, and **show routing-instances** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 55.1.12.1/24;
    }
    family mpls {
      maximum-labels 5;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 55.1.13.1/24;
    }
    family mpls {
      maximum-labels 5;
    }
  }
}
```



```

ge-0/0/5 {
  unit 0 {
    family inet {
      address 55.1.17.1/24;
    }
  }
}

```

user@PE1# show routing-options

```

autonomous-system 65000;
forwarding-table {
export load-balance-policy;
  chained-composite-next-hop {
    ingress {
      l3vpn;
    }
  }
}

```

user@PE1# show protocols

```

mpls {
  interface ge-0/0/0.0;
  interface ge-0/0/1.0;
  label-range {
    static-label-range 1000000 1000999;
  }
}
bgp {
  group pe {
    type internal;
    local-address 128.9.147.211;
    family inet-vpn {
      unicast;
    }
    neighbor 128.9.146.181;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface lo0.0;
    interface ge-0/0/1.0;
  }
}

```



```

    }
}
bfd {
}
source-packet-routing {
    segment-list sl-15-primary {
        hop-1 ip-address 55.1.13.3;
        hop-2 label 1000134;
        hop-3 label 1000145;
    }
    segment-list sl-15-backup {
        hop-1 ip-address 55.1.12.2;
        hop-2 label 1000123;
        hop-3 label 1000134;
        hop-4 label 1000145;
    }
    source-routing-path lsp-15 {
        to 128.9.146.181;
        binding-sid 1000999;
        primary {
            sl-15-primary;
        }
        secondary {
            sl-15-backup;
        }
    }
}
}

```

```

user@PE1# show policy-options
policy-statement VPN-A-export {
    term a {
        from protocol [ ospf direct ];
        then {
            community add VPN-A;
            accept;
        }
    }
    term b {
        then reject;
    }
}
policy-statement VPN-A-import {
    term a {
        from {

```



```

        protocol bgp;
        community VPN-A;
    }
    then accept;
}
term b {
    then reject;
}
}
policy-statement bgp-to-ospf {
    from {
        protocol bgp;
        route-filter 55.1.0.0/16 orlonger;
    }
    then accept;
}
policy-statement load-balance-policy {
    then {
        load-balance per-packet;
    }
}
community VPN-A members target:65000:1;

```

```

user@PE1# show routing-instances
VRF1 {
    instance-type vrf;
    interface ge-0/0/5.0;
    route-distinguisher 128.9.147.211:1;
    vrf-import VPN-A-import;
    vrf-export VPN-A-export;
    vrf-table-label;
    protocols {
        ospf {
            export bgp-to-ospf;
            area 0.0.0.0 {
                interface ge-0/0/5.0;
            }
        }
    }
}

```

Configuring Device PE2

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the interfaces.

```
[edit interfaces]
set ge-0/0/0 unit 0 family inet address 10.10.12.2/24
set ge-0/0/0 unit 0 family mpls

set ge-0/0/1 unit 0 family inet address 10.10.23.2/24
set ge-0/0/1 unit 0 family mpls
```

2. Configure the static LSP for protocol MPLS.

```
[edit protocols mpls static-label-switched-path]
set adj-23 segment 1000123
set adj-23 segment next-hop 10.10.23.3
set adj-23 segment pop
set adj-21 segment 1000221
set adj-21 segment next-hop 10.10.12.1
set adj-21 segment pop
```

3. Configure interfaces and static label range for protocol MPLS.

```
[edit protocols mpls]
set interface ge-0/0/0.0
set interface ge-0/0/1.0
set label-range static-label-range 1000000 1000999
```

4. Configure interfaces for protocol OSPF.

```
[edit protocols ospf area 0.0.0.0]
set interface ge-0/0/0.0
set interface ge-0/0/1.0
```

Results

From configuration mode on router PE2, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

user@PE2# **show interfaces**

```
ge-0/0/0 {
  unit 0 {
    family inet {
      address 55.1.12.2/24;
    }
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 55.1.23.2/24;
    }
    family mpls;
  }
}
```

user@PE2# **show protocols**

```
mpls {
  static-label-switched-path adj-23 {
    segment {
      1000123;
      next-hop 55.1.23.3;
      pop;
    }
  }
  static-label-switched-path adj-21 {
    segment {
      1000221;
      next-hop 55.1.12.1;
      pop;
    }
  }
  interface ge-0/0/0.0;
  interface ge-0/0/1.0;
  label-range {
    static-label-range 1000000 1000999;
  }
}
ospf {
  area 0.0.0.0 {
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
```



```
}
}
```

Verification

IN THIS SECTION

- Verifying Route Entry of Routing Table inet.3 of Router PE1 | 1850
- Verifying Route Table Entries of Routing Table mpls.0 of Router PE1 | 1851
- Verifying SPRING Traffic Engineered LSP of Router PE1 | 1851
- Verifying SPRING Traffic Engineered LSPs on the Ingress Router of Router PE1 | 1852
- Verifying the Routing Table Entries of Routing Table mpls.0 of Router PE2 | 1853
- Verifying the Status of Static MPLS LSP Segments of Router PE2 | 1854

Confirm that the configuration is working properly.

Verifying Route Entry of Routing Table inet.3 of Router PE1

Purpose

Verify the route entry of routing table inet.3 of router PE1.

Action

From operational mode, enter the **show route table inet.3** command.

```
user@PE1> show route table inet.3
```

```
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.146.181/32    *[SPRING-TE/8] 03:09:26, metric 1
                    > to 10.10.13.3 via ge-0/0/1.0, Push 1000145, Push 1000134(top)

                    to 10.10.12.2 via ge-0/0/0.0, Push 1000145, Push 1000134,
                    Push 1000123(top)
```

Meaning

The output displays the ingress routes of segment routing tunnels.

Verifying Route Table Entries of Routing Table mpls.0 of Router PE1

Purpose

Verify the route entries of routing table mpls.0

Action

From operational mode, enter the **show route table mpls.0** command.

user@PE1> **show route table mpls.0**

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 03:25:52, metric 1
            Receive
1          *[MPLS/0] 03:25:52, metric 1
            Receive
2          *[MPLS/0] 03:25:52, metric 1
            Receive
13         *[MPLS/0] 03:25:52, metric 1
            Receive
16         *[VPN/0] 03:25:52
            > via lsi.0 (VRF1), Pop
1000999    *[SPRING-TE/8] 03:04:03, metric 1
            > to 10.10.13.3 via ge-0/0/1.0, Swap 1000145, Push 1000134(top)

            to 10.10.12.2 via ge-0/0/0.0, Swap 1000145, Push 1000134,
            Push 1000123(top)
```

Meaning

The output displays the SID labels of segment routing tunnels.

Verifying SPRING Traffic Engineered LSP of Router PE1

Purpose

Verify SPRING traffic engineered LSPs on the ingress routers.

Action

From operational mode, enter the **show spring-traffic-engineering overview** command.

user@PE1> **show spring-traffic-engineering overview**


```

Overview of SPRING-TE:
  Route preference: 8
  Number of LSPs: 1 (Up: 1, Down: 0)
  External controllers:
    < Not configured >

```

Meaning

The output displays the overview of SPRING traffic engineered LSPs on the ingress router.

Verifying SPRING Traffic Engineered LSPs on the Ingress Router of Router PE1

Purpose

Verify SPRING traffic engineered LSPs on the ingress router.

Action

From operational mode, enter the **show spring-traffic-engineering lsp detail** command.

```
user@PE1# show spring-traffic-engineering lsp detail
```

```

Name: lsp-15
To: 192.168.146.181
State: Up
  Path: sl-15-primary
  Outgoing interface: ge-0/0/1.0
  BFD status: N/A (Up: 0, Down: 0)
  SR-ERO hop count: 3
    Hop 1 (Strict):
      NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.10.13.3
      SID type: None
    Hop 2 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 1000134
    Hop 3 (Strict):
      NAI: None
      SID type: 20-bit label, Value: 1000145
  Path: sl-15-backup
  Outgoing interface: ge-0/0/0.0
  BFD status: N/A (Up: 0, Down: 0)
  SR-ERO hop count: 4
    Hop 1 (Strict):
      NAI: IPv4 Adjacency ID, 0.0.0.0 -> 10.10.12.2
      SID type: None

```



```

Hop 2 (Strict):
  NAI: None
  SID type: 20-bit label, Value: 1000123
Hop 3 (Strict):
  NAI: None
  SID type: 20-bit label, Value: 1000134
Hop 4 (Strict):
  NAI: None
  SID type: 20-bit label, Value: 1000145

Total displayed LSPs: 1 (Up: 1, Down: 0)

```

Meaning

The output displays details of SPRING traffic engineered LSPs on the ingress router

Verifying the Routing Table Entries of Routing Table mpls.0 of Router PE2

Purpose

Verify the routing table entries of routing table mpls.0 of router PE2.

Action

From operational mode, enter the **show route table mpls.0** command.

```
user@PE2> 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] 03:22:29, metric 1
            Receive
1          *[MPLS/0] 03:22:29, metric 1
            Receive
2          *[MPLS/0] 03:22:29, metric 1
            Receive
13         *[MPLS/0] 03:22:29, metric 1
            Receive
1000123    *[MPLS/6] 03:22:29, metric 1
            > to 10.10.23.3 via ge-0/0/1.0, Pop
1000123(S=0) *[MPLS/6] 03:22:29, metric 1
            > to 10.10.23.3 via ge-0/0/1.0, Pop
1000221    *[MPLS/6] 03:22:29, metric 1
            > to 10.10.12.1 via ge-0/0/0.0, Pop

```



```
1000221(S=0)          *[MPLS/6] 03:22:29, metric 1
                      > to 10.10.12.1 via ge-0/0/0.0, Pop
```

Verifying the Status of Static MPLS LSP Segments of Router PE2

Purpose

Verify the status of MPLS LSP segments of router PE2.

Action

From operational mode, enter the **show mpls static-lsp** command.

```
user@PE2> show mpls static-lsp
```

```
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0

Segment LSPs:
LSPname                SID-label    State
adj-21                 1000221      Up
adj-23                 1000123      Up
Total 2, displayed 2, Up 2, Down 0
```

Meaning

The output displays the status of static MPLS LSP segments of router PE2.

Release History Table

Release	Description
19.4R1	You can configure a tunnel template for PCE-initiated segment routing LSPs to pass down two additional parameters for these LSPs - Bidirectional forwarding detection (BFD) and LDP tunneling.
19.1R1	Starting in Junos OS Release 19.1R1, a commit check feature is introduced to ensure that all the segment lists contributing to the colored routes have the minimum label present for all hops.
19.1R1	Starting in Junos OS Release 19.1R1, this requirement does not apply, as the first hop of the non-colored static LSPs now provides support for SID labels, in addition to IP addresses. With the first hop label support, MPLS fast reroute (FRR) and weighted equal-cost multipath is enabled for resolving the static non-colored segment routing LSPs, similar to colored static LSPs.
18.2R1	Starting in Junos OS Release 18.2R1, statically configured non-colored segment routing LSPs on the ingress device are reported to the Path Computation Element (PCE) through a Path Computation Element Protocol (PCEP) session.

RELATED DOCUMENTATION

[segment](#) | [2768](#)
[segment-list](#) | [2770](#)
[source-routing-path](#) | [2782](#)

9

PART

Troubleshooting MPLS

Troubleshooting MPLS | **1859**

Troubleshooting MPLS

IN THIS CHAPTER

- Verify MPLS Interfaces | 1866
- Verify the MPLS Configuration | 1868
- Checklist for Checking the MPLS Layer | 1871
- Checking the MPLS Layer | 1872
- Verify That Node-Link Protection Is Up | 1893
- Verify That Link Protection Is Up | 1900
- Many-to-One Link Protection (Facility Backup) Overview | 1905
- Verify One-to-One Backup | 1907
- Verify That the Primary Path Is Operational | 1914
- Verify That the Secondary Path Is Established | 1916
- Verify the LSP | 1919
- Verify the LSP Route on the Transit Router | 1923
- Verify the LSP Route on the Ingress Router | 1924
- Verify MPLS Labels with the traceroute Command | 1926
- Verify MPLS Labels with the ping Command | 1928
- Take Appropriate Action | 1929
- Verify the LSP Again | 1931
- Checklist for Working with the Layered MPLS Troubleshooting Model | 1935
- Understanding the Layered MPLS Troubleshooting Model | 1935
- Checklist for Verifying the Physical Layer | 1943
- Verifying the Physical Layer | 1944
- Verify the LSP | 1953
- Verify Router Connection | 1955
- Verify Interfaces | 1956
- Take Appropriate Action | 1957
- Verify the LSP Again | 1958
- Checklist for Checking the Data Link Layer | 1960
- Checking the Data Link Layer | 1961

- [Verify the LSP | 1975](#)
- [Verify Interfaces | 1977](#)
- [Take Appropriate Action | 1981](#)
- [Verify the LSP Again | 1983](#)
- [Checklist for Verifying the IP and IGP Layers | 1988](#)
- [Verifying the IP and IGP Layers | 1990](#)
- [Verifying the IP Layer | 1992](#)
- [Verify the LSP | 2006](#)
- [Verify IP Addressing | 2007](#)
- [Verify Neighbors or Adjacencies at the IP Layer | 2009](#)
- [Take Appropriate Action | 2014](#)
- [Verify the LSP Again | 2015](#)
- [Verifying the OSPF Protocol | 2019](#)
- [Verify the LSP | 2035](#)
- [Verify OSPF Interfaces | 2039](#)
- [Verify OSPF Neighbors | 2041](#)
- [Verify the LSP Again | 2042](#)
- [Verify the LSP | 2046](#)
- [Verify IS-IS Adjacencies and Interfaces | 2048](#)
- [Verify the IS-IS Configuration | 2050](#)
- [Verify the LSP Again | 2051](#)
- [Checklist for Checking the RSVP Layer | 2055](#)
- [Checking the RSVP Layer | 2056](#)
- [Verify the LSP | 2073](#)
- [Verify RSVP Sessions | 2075](#)
- [Verify RSVP Neighbors | 2077](#)
- [Verify RSVP Interfaces | 2079](#)
- [Verify the RSVP Protocol Configuration | 2081](#)
- [Take Appropriate Action | 2082](#)
- [Verify the LSP Again | 2083](#)
- [Checklist for Determining LSP Status | 2087](#)
- [Determining LSP Statistics | 2088](#)
- [Checklist for Verifying LSP Use | 2090](#)
- [Verifying LSP Use in Your Network | 2090](#)

- Verifying an LSP on the Ingress Router | 2094
- Verifying an LSP on a Transit Router | 2096
- Verify That Load Balancing Is Working | 2097
- Example: Load-Balanced MPLS Network | 2101
- Router Configurations for the Load-Balanced MPLS Network | 2102
- Traffic Flows Before Load Balancing | 2117
- Verify the Operation of Uneven Bandwidth Load Balancing | 2119
- Checklist for Collecting Crash Data | 2121
- Understand Crash Data Collection | 2123
- Collect Crash Data for a Routing Engine Kernel | 2124
- Check the Routing Engine Core Files | 2128
- List the Core Files | 2129
- Compress the vmcore File | 2130
- Log Software Version Information | 2131
- Open a Case with JTAC | 2132
- Collect Crash Data for Routing Engine Daemons | 2132
- Collect and Send Routing Engine Crash Data to JTAC | 2137
- Check for Daemon Core Files | 2138
- List the Daemon Core Files | 2139
- Compress the Daemon Core Files | 2141
- Collect Crash Data for the Packet Forwarding Engine Microkernel | 2141
- Display the Crash Stack Traceback and Registration Information | 2151
- Clear the NVRAM Contents | 2156
- Check Packet Forwarding Engine Microkernel Core Files | 2157
- List the Core Files Generated by the Crash | 2157
- Compress the Core Files | 2158
- Configure a Primary Path | 2159
- Ensuring That Secondary Paths Establish When Resources Are Diminished | 2161
- One-to-One Backup Overview | 2163
- Configure Link Protection | 2164
- Configuring and Verifying Link Protection | 2166
- Configure Node-Link Protection | 2173
- Configuring and Verifying Node-Link Protection | 2175
- Configure IS-IS as the IGP | 2185

- Verify That IS-IS Adjacencies Are Established | 2194
- Configure OSPF as the IGP | 2196
- Set Up BGP on Routers in Your Network | 2201
- Define the Local Autonomous System | 2209
- Enable MPLS and RSVP | 2211
- Enable MPLS and RSVP on Routers | 2214
- Enable MPLS on Transit Router Interfaces | 2216
- Verifying the MPLS Configuration | 2217
- Verify the RSVP Protocol | 2226
- Define a Load-Balancing Policy | 2228
- Use the traceroute Command to Verify MPLS Labels | 2229
- Apply the Load-Balancing Policy to the Forwarding Table | 2230
- Fast Reroute Problem Overview | 2232
- Problem Establishing a GRE Tunnel Checklist | 2257
- Troubleshooting GMPLS and GRE Tunnel | 2258
- Verify Protocol Families | 2281
- Determining LSP Status | 2285
- Check the Status of the LSP | 2291
- Display Extensive Status About the LSP | 2292
- Checking That RSVP Path Messages Are Sent and Received | 2297
- Determining the Current RSVP Neighbor State | 2299
- Take Appropriate Action | 2300
- Examine BGP Routes | 2302
- CLI Operational Mode Top-Level Commands | 2303
- CLI Keyboard Shortcuts | 2306
- Manage Output at the ---(more)--- Prompt | 2307
- Working with Problems on Your Network | 2308
- Isolating a Broken Network Connection | 2309
- Display Junos OS Information | 2310
- Display Version Information for Junos OS Packages | 2312
- Display the Current Active Router Configuration | 2313
- Copy Junos OS to the Router | 2318
- Add New Software | 2319
- Compare Information Logged Before and After the Upgrade | 2320

- [Displaying LSP Status Events | 2320](#)
- [Call Was Cleared by RSVP Event | 2324](#)
- [Change in Active Path Event | 2324](#)
- [Clear Call Event | 2325](#)
- [Deselected as Active Event | 2326](#)
- [Link Protection Down Event | 2327](#)
- [Originate Call Event | 2329](#)
- [ResvTear Received Event | 2329](#)
- [Session Preempted Event | 2330](#)
- [Displaying General LSP Error Events | 2331](#)
- [Admission Control Failure Event | 2333](#)
- [Explicit Route: Bad Loose Route Event | 2333](#)
- [Explicit Route: Bad Strict Route Event | 2335](#)
- [Explicit Route: Format Error Event | 2337](#)
- [Explicit Route: Wrong Delivery Event | 2338](#)
- [Invalid Destination Address Event | 2340](#)
- [Invalid Filter for Policing Event | 2341](#)
- [MPLS Graceful Restart: Recovery Failed Event | 2341](#)
- [MPLS Label Allocation Failure Event | 2342](#)
- [Non-RSVP Capable Router Detected Event | 2343](#)
- [No Route Toward Destination Event | 2344](#)
- [Unsupported Traffic Class Event | 2345](#)
- [CSPF: Computation Result Accepted Event | 2345](#)
- [CSPF: Reroute Due to Re-Optimization Event | 2346](#)
- [Retry Limit Exceeded Event | 2347](#)
- [Log the Software Version Information | 2348](#)
- [Log the Hardware Version Information | 2349](#)
- [Log the System Boot-Message Information | 2351](#)
- [Log the BGP, IS-IS, and OSPF Adjacency Information | 2353](#)
- [Back Up the Currently Running and Active File System | 2355](#)
- [Reinstall Junos OS | 2355](#)
- [Reconfigure Junos OS | 2356](#)
- [Configure Host Names, Domain Names, and IP Addresses | 2360](#)
- [Check Network Connectivity | 2361](#)

- Automatic Autobandwidth Adjustment Failed Event | 2361
- Configuring Automatic Bandwidth Allocation for LSPs | 2363
- Displaying DiffServ-Aware Traffic-Engineered LSP Events | 2371
- Unsupported Traffic Class Event | 2372
- Traffic Class Value Out of Allowed Range Event | 2373
- The Combination of Setup Priority and Traffic Class Is Not One of the Configured TE Classes Event | 2373
- RSVP Error, Subcode 7, Signal Type Does Not Match Link Encoding Event | 2374
- Unacceptable Label Value Event | 2374
- Unsupported Switching Type Event | 2375
- Gather Component Alarm Information | 2375
- Case Study for a CSPF Failure | 2377
- Examining a CSPF Failure | 2385
- Verify the CSPF Failure | 2396
- Examining the Hello Message | 2398
- Displaying the Status of IS-IS Adjacencies | 2401
- Check OSPF on a Stub Router | 2405
- Checklist for Verifying the BGP Protocol and Peers | 2406
- Verify BGP Peers | 2407
- Examine BGP Routes and Route Selection | 2418
- Examine the Local Preference Selection | 2426
- Examine the Multiple Exit Discriminator Route Selection | 2427
- Examine the EBGp over IBGP Selection | 2429
- Examine the IGP Cost Selection | 2430
- Examine Routes in the Forwarding Table | 2432
- Ping the Egress Router | 2433
- View the RSVP Log File on Transit Routers | 2434
- Check the RSVP Log File on the Egress Router | 2436
- Determine and Correct the Problem on the Egress Router | 2437
- Check the Routing CPU Memory Usage | 2439
- Run Snmpwalk from an NMS System to a Juniper Router | 2449
- Configure Trace Operations for SNMP | 2451
- Query a MIB With SNMPGet | 2452
- Check CPU Utilization | 2453
- Check CPU Utilization per Process | 2455

- Retrieve Version Information about Router Software Components | 2459
- Checklist for Displaying Basic Chassis Information | 2460
- Display Basic Chassis Information | 2461
- Maintain a Single Configuration File for Both Routing Engines | 2464
- Configure the New Group | 2468
- Apply the New Group | 2470
- Viewing Files and Directories on a Device Running Junos OS | 2472
- Display File Contents | 2477
- Rename a File on a Router | 2478
- Delete a File on a Router | 2479
- Check the Time on a Router | 2479
- Check for Users in Configuration Mode | 2480
- Check the Commands That Users Are Entering | 2481
- Configure the Log File for Tracking CLI Commands | 2484
- Check When the Last Configuration Change Occurred | 2485
- Configure Configuration Change Tracking | 2487
- Display a Log File | 2488
- Configure IS-IS-Specific Options | 2489
- Displaying Detailed IS-IS Protocol Information | 2498
- Analyzing IS-IS Link-State PDUs in Detail | 2501
- Configure OSPF-Specific Options | 2504
- Diagnose OSPF Session Establishment Problems | 2511
- Analyze OSPF Link-State Advertisement Packets in Detail | 2516
- Chassis | 2517
- Physical Interface Cards | 2518
- Routing Engine | 2518
- Compare Information Logged Before and After the Reinstall | 2519
- Back Up the New Software | 2519
- Monitor Hardware Components | 2520
- Log Software Version Information | 2520
- Hardware Components | 2521

Verify MPLS Interfaces

Purpose

If the MPLS protocol is not configured correctly on the routers in your network, the interfaces are not able to perform MPLS switching.

NOTE: For a labeled route to be resolved over an interface, **family mpls** must be configured at the **[edit interfaces]** hierarchy level for the route to be successfully resolved. When the interface is not configured with **family mpls**, labelled routes do not get resolved.

Action

To verify MPLS interfaces, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show mpls interface
```

Sample Output 1

The following sample output is for all routers in the network shown in [“MPLS Network Topology” on page 2217](#).

```
user@R1> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
```

```
user@R2> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>
so-0/0/3.0     Up         <none>
```

```
user@R3> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
```



```
so-0/0/1.0      Up      <none>
so-0/0/2.0      Up      <none>
so-0/0/3.0      Up      <none>
```

user@R4> **show mpls interface**

Interface	State	Administrative groups
so-0/0/0.0	Up	<none>
so-0/0/1.0	Up	<none>
so-0/0/2.0	Up	<none>
so-0/0/3.0	Up	<none>

user@R5> **show mpls interface**

Interface	State	Administrative groups
so-0/0/0.0	Up	<none>
so-0/0/1.0	Up	<none>
so-0/0/2.0	Up	<none>

user@R6> **show mpls interface**

Interface	State	Administrative groups
so-0/0/0.0	Up	<none>
so-0/0/1.0	Up	<none>
so-0/0/2.0	Up	<none>
so-0/0/3.0	Up	<none>

Sample Output 2

user@R6> **show mpls interface**

Interface	State	Administrative groups
so-0/0/0.0	Up	<none>
so-0/0/1.0	Up	<none>
so-0/0/3.0	Up	<none>

so-0/0/2.0 is missing

Sample Output 3

```
user@host> show mpls interface
MPLS not configured
```

Meaning

Sample Output 1 shows that all MPLS interfaces on all routers in the network are enabled (**Up**) and can perform MPLS switching. If you fail to configure the correct interface at the **[edit protocols mpls]** hierarchy level or include the **family mpls** statement at the **[edit interfaces type-fpc/pic/port unit number]** hierarchy level, the interface cannot perform MPLS switching, and does not appear in the output for the **show mpls interface** command.

Administrative groups are not configured on any of the interfaces shown in the example network in “[MPLS Network Topology](#)” on page 2217. However, if they were, the output would indicate which affinity class bits are enabled on the router.

Sample Output 2 shows that interface **so-0/0/2.0** is missing and therefore might be incorrectly configured. For example, the interface might not be included at the **[edit protocols mpls]** hierarchy level, or the **family mpls** statement might not be included at the **[edit interfaces type-fpc/pic/port unit number]** hierarchy level. If the interface is configured correctly, RSVP might not have signaled over this interface yet. For more information on determining which interface is incorrectly configured, see “[Verify Protocol Families](#)” on page 2222.

Sample Output 3 shows that the MPLS protocol is not configured at the **[edit protocols mpls]** hierarchy level.

Verify the MPLS Configuration

Purpose

After you have checked the transit and ingress routers, use the **traceroute** command to verify the BGP next hop, and used the **ping** command to verify the active path, you can check for problems with the MPLS configuration at the **[edit protocols mpls]** and **[edit interfaces]** hierarchy levels.

NOTE: For a labeled route to be resolved over an interface, **family mpls** must be configured at the **[edit interfaces]** hierarchy level for the route to be successfully resolved. When the interface is not configured with **family mpls**, labelled routes do not get resolved.

Action

To verify the MPLS configuration, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show configuration protocols mpls
user@host> show configuration interfaces
```

Sample Output 1

```
user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}

user@R3> show configuration protocols mpls
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;

user@R6> show configuration protocols mpls
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured
```


Sample Output 2

```
user@R6> show configuration interfaces
so-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.56.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.46.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.26.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/3 {
    unit 0 {
        family inet {
            address 10.1.36.2/30;
        }
        family iso;
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.148/21;
        }
    }
}
```



```
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.6/32;
        address 127.0.0.1/32;
      }
      family iso {
        address 49.0003.1000.0000.0006.00;
      }
    }
  }
}
```

Meaning

Sample Output 1 from the ingress, transit, and egress routers shows that the configuration of interfaces on egress router **R6** is incorrect. Interface **so-0/0/3.0** is included as inactive at the **[edit protocols mpls]** hierarchy level, when it should be active because it is the interface through which the LSP travels.

Sample Output 2 shows that interfaces are correctly configured for MPLS on egress router **R6**. The interfaces are also correctly configured on the ingress and transit routers (not shown).

Checklist for Checking the MPLS Layer

Problem

Description: This checklist provides the steps and commands for checking the Multiprotocol Label Switching (MPLS) layer of the layered MPLS model. The checklist provides links to an overview of the MPLS layer and more detailed information about the commands used to investigate the problem.

[Table 50 on page 1871](#) provides commands for checking the MPLS layer.

Solution

Table 50: Checklist for Checking the MPLS Layer

Tasks	Command or Action
“Checking the MPLS Layer” on page 1872	
1. Verify the LSP on page 1874	show mpls lsp show mpls lsp extensive show mpls lsp name <i>name</i> show mpls lsp name <i>name</i> extensive

Table 50: Checklist for Checking the MPLS Layer (*continued*)

Tasks	Command or Action
2. Verify the LSP Route on the Transit Router on page 1878	<code>show route table mpls.0</code>
3. Verify the LSP Route on the Ingress Router on page 1880	<code>show route <i>destination</i></code>
4. Verify MPLS Labels with the traceroute Command on page 1882	<code>traceroute <i>hostname</i></code>
5. Verify MPLS Labels with the ping Command on page 1883	On the ingress router: <code>ping mpls rsvp <i>lsp-name</i> detail</code>
6. Verify the MPLS Configuration on page 1868	<code>show configuration protocols mpls</code> <code>show configuration interfaces</code>
7. Take Appropriate Action on page 1888	The following sequence of commands addresses the specific problem described in this topic: <code>edit</code> <code>edit protocols mpls</code> <code>[edit protocols mpls]</code> <code>show</code> <code>activate interface so-0/0/3.0</code> <code>show</code> <code>commit</code>
8. Verify the LSP Again on page 1889	<code>show mpls lsp extensive</code>

Checking the MPLS Layer

Purpose

After you have configured the label-switched path (LSP), issued the `show mpls lsp` command, and determined that there is an error, you might find that the error is not in the physical, data link, Internet Protocol (IP), interior gateway protocol (IGP), or Resource Reservation Protocol (RSVP) layers. Continue investigating the problem at the MPLS layer of the network.

[Figure 145 on page 1873](#) illustrates the MPLS layer of the layered MPLS model.

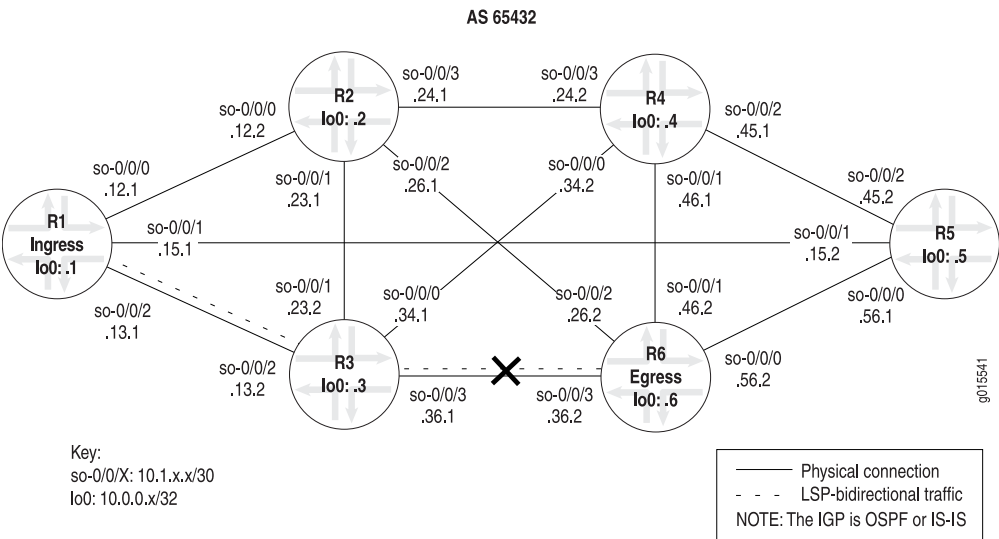
Figure 145: Checking the MPLS Layer

BGP Layer	traceroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
↩ IGP and IP Layers Functioning ↪	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive "JUNOS Interfaces Operations Guide"
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

With the MPLS layer, you check whether the LSP is up and functioning correctly. If the network is not functioning at this layer, the LSP does not work as configured.

Figure 146 on page 1873 illustrates the MPLS network used in this topic.

Figure 146: MPLS Network Broken at the MPLS Layer



The network shown in [Figure 146 on page 1873](#) is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, the reverse LSP is down without a path from **R6** to **R1**.

The cross shown in [Figure 146 on page 1873](#) indicates where the LSP is broken. Some possible reasons the LSP is broken might include an incorrectly configured MPLS protocol, or interfaces that are incorrectly configured for MPLS.

In the network shown in [Figure 146 on page 1873](#), a configuration error on egress router **R6** prevents the LSP from traversing the network as expected.

To check the MPLS layer, follow these steps:

1. [Verify the LSP | 1874](#)
2. [Verify the LSP Route on the Transit Router | 1878](#)
3. [Verify the LSP Route on the Ingress Router | 1880](#)
4. [Verify MPLS Labels with the traceroute Command | 1882](#)
5. [Verify MPLS Labels with the ping Command | 1883](#)
6. [Verify the MPLS Configuration | 1885](#)
7. [Take Appropriate Action | 1888](#)
8. [Verify the LSP Again | 1889](#)

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To verify that the LSP is up, enter some or all of the following commands from the ingress router:

```
user@host> show mpls lsp
user@host> show mpls lsp extensive
user@host> show mpls lsp name name
user@host> show mpls lsp name name extensive
```


Sample Output 1

```

user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6     10.0.0.1      Dn      0  -              R1-to-R6
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.1     10.0.0.6      Dn      0  -              R6-to-R1
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```


Sample Output 2

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 22 second(s).
    1 Nov  2 14:43:38  CSPF failed: no route toward 10.0.0.6 [175 times]
  Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 13 second(s).
    1 Nov  2 14:38:12  CSPF failed: no route toward 10.0.0.1 [177 times]

```



```
Created: Tue Nov  2 13:12:22 2004
Total 1 displayed, Up 0, Down 1
```

```
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 3

```
user@R1> show mpls lsp name R1-to-R6
Ingress LSP: 1 sessions

```

To	From	State	Rt	ActivePath	P	LSPname
10.0.0.6	10.0.0.1	Dn	0	-		R1-to-R6

```
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 4

```
user@R1> show mpls lsp name R1-to-R6 extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 10 second(s).
    1 Nov  2 14:51:53 CSPF failed: no route toward 10.0.0.6[192 times]
```



```

Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 shows a brief description of the state of the LSP for the ingress, transit, and egress routers. Output from ingress router **R1** and egress router **R6** shows that both LSPs are down, **R1-to-R6** and **R6-to-R1**. With the configured LSPs on **R1** and **R6**, we would expect egress LSP sessions on both **R1** and **R6**. In addition, transit router **R3** has no transit sessions.

Sample Output 2 shows all information about the LSPs, including all past state history and the reason why an LSP failed. Output from **R1** and **R6** indicates that there is no route to the destination because the Constrained Shortest Path First (CSPF) algorithm failed.

Sample Outputs 3 and 4 show examples of the output for the **show mpls lsp name** command with the **extensive** option. In this instance, the output is very similar to the **show mpls lsp** command because only one LSP is configured in the example network in [Figure 146 on page 1873](#). However, in a large network with many LSPs configured, the results would be quite different between the two commands.

Verify the LSP Route on the Transit Router

Purpose

If the LSP is up, the LSP route should appear in the **mpls.0** routing table. MPLS maintains an MPLS path routing table (**mpls.0**), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP. If routes are not present in the output for the transit router, check the MPLS protocol configuration on the ingress and egress routers.

Action

To verify the LSP route on the transit router, enter the following command from the transit router:

```
user@host> show route table mpls.0
```

Sample Output 1

```
user@R3> show route table mpls.0
```



```

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 16w2d 21:52:40, metric 1
            Receive
1          *[MPLS/0] 16w2d 21:52:40, metric 1
            Receive
2          *[MPLS/0] 16w2d 21:52:40, metric 1
            Receive

```

Sample Output 2

user@R3> show route table mpls.0

```

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 16w2d 22:26:08, metric 1
            Receive
1          *[MPLS/0] 16w2d 22:26:08, metric 1
            Receive
2          *[MPLS/0] 16w2d 22:26:08, metric 1
            Receive
100864     *[RSVP/7] 00:07:23, metric 1
            > via so-0/0/2.0, label-switched-path R6-to-R1
100864(S=0) *[RSVP/7] 00:07:23, metric 1
            > via so-0/0/2.0, label-switched-path R6-to-R1
100880     *[RSVP/7] 00:07:01, metric 1
            > via so-0/0/3.0, label-switched-path R1-to-R6
100880(S=0) *[RSVP/7] 00:07:01, metric 1
            > via so-0/0/3.0, label-switched-path R1-to-R6

```

Meaning

Sample Output 1 from transit router **R3** shows three route entries in the form of MPLS label entries. These MPLS labels are reserved MPLS labels defined in RFC 3032, and are always present in the **mpls.0** routing table, regardless of the state of the LSP. The incoming labels assigned by RSVP to the upstream neighbor are missing from the output, indicating that the LSP is down. For more information on MPLS label entries, see [“Checklist for Verifying LSP Use” on page 2090](#).

In contrast, Sample Output 2 shows the MPLS labels and routes for a correctly configured LSP. The three reserved MPLS labels are present, and the four other entries represent the incoming labels assigned by RSVP to the upstream neighbor. These four entries represent two routes. There are two entries per route

because the stack values in the MPLS header may be different. For each route, the second entry **100864 (S=0)** and **100880 (S=0)** indicates that the stack depth is not 1, and additional label values are included in the packet. In contrast, the first entry, **100864** and **100880** has an inferred S=1 value which indicates a stack depth of 1 and makes each label the last label in that particular packet. The dual entries indicate that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

Verify the LSP Route on the Ingress Router

Purpose

Check whether the LSP route is included in the active entries in the **inet.3** routing table for the specified address.

Action

To verify the LSP route, enter the following command from the ingress router:

```
user@host> show route destination
```

Sample Output 1

```
user@R1> show route 10.0.0.6
```

```
inet.0 : 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.6/32      *[IS-IS/18] 6d 01:41:37, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
                  to 10.1.13.2 via so-0/0/2.0
```

```
user@R6> show route 10.0.0.1
```

```
inet.0 : 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.1/32      *[IS-IS/18] 5d 01:01:38, metric 20
                  to 10.1.56.1 via so-0/0/0.0
                  > to 10.1.26.1 via so-0/0/2.0
                  to 10.1.36.1 via so-0/0/3.0
```


Sample Output 2

user@R1> show route 10.0.0.6

```
inet.0: 28 destinations, 28 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.6/32      *[IS-IS/18] 6d 02:13:42, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
                  to 10.1.13.2 via so-0/0/2.0
```

```
inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.6/32      *[RSVP/7] 00:08:07, metric 20
                  > via so-0/0/2.0, label-switched-path R1-to-R6
```

user@R6> show route 10.0.0.1

```
inet.0: 29 destinations, 29 routes (28 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.1/32      *[IS-IS/18] 5d 01:34:03, metric 20
                  to 10.1.56.1 via so-0/0/0.0
                  > to 10.1.26.1 via so-0/0/2.0
                  to 10.1.36.1 via so-0/0/3.0
```

```
inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
10.0.0.1/32      *[RSVP/7] 00:10:39, metric 20
                  > via so-0/0/3.0, label-switched-path R6-to-R1
```

Meaning

Sample Output 1 shows entries in the **inet.0** routing table only. The **inet.3** routing table is missing from the output because the LSP is not working. The **inet.0** routing table is used by interior gateway protocols (IGPs) and Border Gateway Protocol (BGP) to store routing information. In this case, the IGP is Intermediate System-to-Intermediate System (IS-IS). For more information on the **inet.0** routing table, see the *Junos MPLS Applications Configuration Guide*.

If the LSP was working, we would expect to see entries that include the LSP in the **inet.3** routing table. The **inet.3** routing table is used on ingress routers to route BGP packets to the destination egress router.

BGP uses the **inet.3** routing table on the ingress router to help resolve next-hop addresses. BGP is configured in the example network shown in [Figure 146 on page 1873](#).

Sample Output 2 shows output you should receive when the LSP is up. The output shows both the **inet.0** and **inet.3** routing tables, indicating that LSPs **R1-to-R6** and **R6-to-R1** are available.

Verify MPLS Labels with the traceroute Command

Purpose

Display the route packets take to a BGP destination where the BGP next hop for that route is the LSP egress address. By default, BGP uses the **inet.0** and the **inet.3** routing tables to resolve the next-hop address. When the next-hop address of the BGP route is not the router ID of the egress router, traffic is mapped to IGP routes, not to the LSP. Use the **traceroute** command as a debugging tool to determine whether the LSP is being used to forward traffic.

Action

To verify MPLS labels, enter the following command from the ingress router:

```
user@host> traceroute hostname
```

Sample Output 1

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.12.2 (10.1.12.2)  0.627 ms  0.561 ms  0.520 ms
 2  10.1.26.2 (10.1.26.2)  0.570 ms !N  0.558 ms !N  4.879 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.630 ms  0.545 ms  0.488 ms
 2  10.1.12.1 (10.1.12.1)  0.551 ms !N  0.557 ms !N  0.526 ms !N
```

Sample Output 2

```
user@R1> traceroute 100.100.6.1
to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
```



```

1  10.1.13.2 (10.1.13.2)  0.866 ms  0.746 ms  0.724 ms
    MPLS Label=100912 CoS=0 TTL=1 S=1
2  10.1.36.2 (10.1.36.2)  0.577 ms !N  0.597 ms !N  0.546 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
1  10.1.36.1 (10.1.36.1)  0.802 ms  0.716 ms  0.688 ms
    MPLS Label=100896 CoS=0 TTL=1 S=1
2  10.1.13.1 (10.1.13.1)  0.570 ms !N  0.568 ms !N  0.546 ms !N

```

Meaning

Sample Output 1 shows that BGP traffic is not using the LSP, consequently MPLS labels do not appear in the output. Instead of using the LSP, BGP traffic is using the IGP (IS-IS, in the example network in [Figure 146 on page 1873](#)) to reach the BGP next-hop LSP egress address. The Junos OS default behavior uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Sample Output 2 is an example of output for a correctly configured LSP. The output shows MPLS labels, indicating that BGP traffic is using the LSP to reach the BGP next hop.

Verify MPLS Labels with the ping Command

Purpose

When you ping a specific LSP, you check that echo requests are sent over the LSP as MPLS packets.

Action

To verify MPLS labels, enter the following command from the ingress router to ping the egress router:

```
user@host> ping mpls rsvp lsp-name detail
```

For example:

```
user@R1> ping mpls rsvp R1-to-R6 detail
```

Sample Output 1

```

user@R1> ping mpls rsvp R1-to-R6 detail
LSP R1-to-R6 - LSP has no active path, exiting.

```



```
user@R6> ping mpls rsvp R6-to-R1 detail
LSP R6-to-R1 - LSP has no active path, exiting.
```

Sample Output 2

```
user@R1> traceroute 10.0.0.6
traceroute to 10.0.0.6 (10.0.0.6), 30 hops max, 40 byte packets
 1  10.1.15.2 (10.1.15.2)  0.708 ms  0.613 ms  0.576 ms
 2  10.0.0.6 (10.0.0.6)  0.763 ms  0.708 ms  0.700 ms

user@R1> ping mpls rsvp R1-to-R6 detail
Request for seq 1, to interface 69, label 100880
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 69, label 100880
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 69, label 100880
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 69, label 100880
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 69, label 100880
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

user@R6> ping mpls rsvp R6-to-R1 detail
Request for seq 1, to interface 70, label 100864
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 70, label 100864
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 70, label 100864
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 70, label 100864
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 70, label 100864
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


Meaning

Sample Output 1 shows that the LSP does not have an active path to forward echo requests, indicating that the LSP is down.

Sample Output 2 is an example of output you should receive when the LSP is up and forwarding packets.

Verify the MPLS Configuration

Purpose

After you have checked the transit and ingress routers, use the **traceroute** command to verify the BGP next hop, and used the **ping** command to verify the active path, you can check for problems with the MPLS configuration at the **[edit protocols mpls]** and **[edit interfaces]** hierarchy levels.

NOTE: For a labeled route to be resolved over an interface, **family mpls** must be configured at the **[edit interfaces]** hierarchy level for the route to be successfully resolved. When the interface is not configured with **family mpls**, labelled routes do not get resolved.

Action

To verify the MPLS configuration, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show configuration protocols mpls
user@host> show configuration interfaces
```

Sample Output 1

```
user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}
```



```

user@R3> show configuration protocols mpls
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;

user@R6> show configuration protocols mpls
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured

```

Sample Output 2

```

user@R6> show configuration interfaces
so-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.56.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.46.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/2 {

```



```

    unit 0 {
        family inet {
            address 10.1.26.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/3 {
    unit 0 {
        family inet {
            address 10.1.36.2/30;
        }
        family iso;
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.148/21;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.6/32;
            address 127.0.0.1/32;
        }
        family iso {
            address 49.0003.1000.0000.0006.00;
        }
    }
}
}

```

Meaning

Sample Output 1 from the ingress, transit, and egress routers shows that the configuration of interfaces on egress router **R6** is incorrect. Interface **so-0/0/3.0** is included as inactive at the **[edit protocols mpls]** hierarchy level, when it should be active because it is the interface through which the LSP travels.

Sample Output 2 shows that interfaces are correctly configured for MPLS on egress router **R6**. The interfaces are also correctly configured on the ingress and transit routers (not shown).

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, an interface is incorrectly configured at the **[edit protocols mpls]** hierarchy level on egress router **R6**.

Solution

To correct the error in this example, follow these steps:

1. Activate the interface in the MPLS protocol configuration on egress router **R6**:

```
user@R6> edit
user@R6# edit protocols mpls
[edit protocols mpls]
user@R6# show
user@R6# activate interface so-0/0/3.0
```

2. Verify and commit the configuration:

```
[edit protocols mpls]
user@R6# show
user@R6# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols mpls

[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured interface

[edit protocols mpls]
```



```

user@R6# activate interface so-0/0/3

[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
interface so-0/0/3.0; <<< Correctly configured interface

[edit protocols mpls]
user@R6# commit
commit complete

```

Meaning

The sample output shows that the incorrectly configured interface **so-0/0/3.0** on egress router **R6** is now activated at the **[edit protocols mpls]** hierarchy level. The LSP can now come up.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the BGP layer has been resolved.

Action

To verify the LSP again, enter the following command from the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
    From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6

```



```

ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
        6 Nov  2 15:48:52 Selected as active path
        5 Nov  2 15:48:52 Record Route:  10.1.13.2 10.1.36.2
        4 Nov  2 15:48:52 Up
        3 Nov  2 15:48:52 Originate Call
        2 Nov  2 15:48:52 CSPF: computation result accepted
        1 Nov  2 15:48:22 CSPF failed: no route toward 10.0.0.6[308 times]
    Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6,  LSPstate: Up , ActiveRoute: 0
LSPname: R6-to-R1 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 159, Since: Tue Nov  2 15:48:30 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39106 protocol 0
PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions

```


Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1

From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
 LSPname: R6-to-R1, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100864, Label out: 3
 Time left: 123, Since: Tue Nov 2 15:35:41 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 39106 protocol 0
 PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.13.1 (so-0/0/2.0) 10 pkts
 RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
 Explct route: 10.1.13.1
 Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
 LSPname: R1-to-R6, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100880, Label out: 3
 Time left: 145, Since: Tue Nov 2 15:36:03 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 48015 protocol 0
 PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.36.2 (so-0/0/3.0) 10 pkts
 RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
 Explct route: 10.1.36.2
 Record route: 10.1.13.1 <self> 10.1.36.2

Total 2 displayed, Up 2, Down 0

user@R6> **show mpls lsp extensive**

Ingress LSP: 1 sessions

10.0.0.1

From: 10.0.0.6, **State: Up, ActiveRoute: 1, LSPname: R6-to-R1**
 ActivePath: (primary)
 LoadBalance: Random


```

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.36.1 10.1.13.1
    6 Nov  2 15:41:44 Selected as active path
    5 Nov  2 15:41:44 Record Route:  10.1.36.1 10.1.13.1
    4 Nov  2 15:41:44 Up
    3 Nov  2 15:41:44 Originate Call
    2 Nov  2 15:41:44 CSPF: computation result accepted
    1 Nov  2 15:41:14 CSPF failed: no route toward 10.0.0.1[306 times]
Created: Tue Nov  2 13:12:21 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
From: 10.0.0.1,  LSPstate: Up,  ActiveRoute: 0
LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left:  157, Since: Tue Nov  2 15:42:06 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 48015 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Verify That Node-Link Protection Is Up

Purpose

After you configure node-link protection, you must check that bypass paths are up. You can also check the number of LSPs protected by the bypass paths. In the network shown in [Figure 7 on page 144](#), two bypass paths should be up: one next-hop bypass path protecting the link between **R1** and **R2** (or next-hop **10.0.12.14**), and a next-next-hop bypass path avoiding **R2**.

Action

To verify node-link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router. You can also issue the commands on transit routers and other routers used in the bypass path for slightly different information.

```
show mpls lsp
show mpls lsp extensive
show rsvp interface
show rsvp interface extensive
show rsvp session detail
```

Sample Output

```
user@R1> show mpls lsp
```

Ingress LSP: 1 sessions

To	From	State	Rt	ActivePath	P	LSPname
192.168.5.1	192.168.1.1	Up	0	via-r2	*	lsp2-r1-to-r5

Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
192.168.1.1	192.168.5.1	Up	0	1 FF	3		- r5-to-r1

Total 1 displayed, **Up 1** , Down 0

Transit LSP: 2 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
----	------	-------	----	-------	---------	----------	---------


```

192.168.0.1      192.168.6.1      Up           0  1 FF 100464      101952  lsp1-r6-to-r0
192.168.6.1      192.168.0.1      Up           0  1 FF  100448           3  r0-to-t6
Total 2 displayed, Up 2, Down 0

```

Meaning

Sample output from **R1** for the **show mpls lsp** command shows a brief description of the state of configured and active LSPs for which **R1** is the ingress, transit, and egress router. All LSPs are up. **R1** is the ingress router for **lsp2-r1-to-r5**, and the egress router for return LSP **r5-to-r1**. Two LSPs transit **R1**, **lsp1-r6-to-r0** and the return LSP **r0-to-t6**. For more detailed information about the LSP, include the **extensive** option when you issue the **show mpls lsp** command.

Sample Output

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up , ActiveRoute: 0, LSPname: lsp2-r1-to-r5
  ActivePath: via-r2 (primary)
    Node/Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-r2 State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.0.12.14(Label=101872) 10.0.24.2(Label=101360) 10.0.45.2(Label=3)
11 Jul 11 14:30:58 Link-protection Up
  10 Jul 11 14:28:28 Selected as active path
  [...Output truncated...]
  Created: Tue Jul 11 14:22:58 2006
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.1.1
  From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r5-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -

```


Recovery label received: -, Recovery label sent: -
 Resv style: 1 FF, Label in: 3, Label out: -
 Time left: 146, Since: Tue Jul 11 14:28:36 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 29228 protocol 0
 PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 362 pkts
 Adspec: received MTU 1500
 PATH sentto: localclient
 RESV rcvfrom: localclient
 Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>
 Total 1 displayed, Up 1, Down 0

Transit LSP: 2 sessions

192.168.0.1

From: 192.168.6.1, **LSPstate: Up**, ActiveRoute: 0
LSPname: lsp1-r6-to-r0, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 101952
 Resv style: 1 SE, Label in: 100464, Label out: 101952
 Time left: 157, Since: Tue Jul 11 14:31:38 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 11131 protocol 0
Node/Link protection desired
Type: Node/Link protected LSP, using Bypass->10.0.12.14->10.0.24.2
1 Jul 11 14:31:38 Node protection up, using Bypass->10.0.12.14->10.0.24.2
 PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 509 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.0.12.14 (fe-0/1/0.0) 356 pkts
 RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts
 Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
 Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1

From: 192.168.0.1, **LSPstate: Up**, ActiveRoute: 0
LSPname: r0-to-t6, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100448, Label out: 3
 Time left: 147, Since: Tue Jul 11 14:31:36 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 23481 protocol 0
 PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts


```

Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.16.2 (so-0/0/3.0) 350 pkts
RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 323 pkts
Explot route: 10.0.16.2
Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning

Sample output from **R1** for the **show mpls lsp extensive** command shows detailed information about all LSPs for which **R1** is the ingress, egress, or transit router, including all past state history and the reason why an LSP failed. All LSPs are up. The main two LSPs **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have node-link protection as indicated by the **Node/Link protection desired** field in the ingress and transit sections of the output. In the ingress section of the output, the **Link-protection Up** field shows that **lsp2-r1-to-r5** has link protection up. In the transit section of the output, the **Type: Node/Link protected LSP** field shows that **lsp1-r6-to-r0** has node-link protection up, and in case of failure will use the bypass LSP **Bypass->10.0.12.14->10.0.24.2**.

Sample Output

```

user@R1> show rsvp interface
RSVP interface: 4 active

```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
fe-0/1/0.0	Up	2	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/1.0	Up	1	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/2.0	Up	0	100%	100Mbps	100Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning

Sample output from **R1** for the **show rsvp interface** command shows four interfaces enabled with RSVP (**Up**). Interface **fe-0/1/0.0** has two active RSVP reservations (**Active resv**) that might indicate sessions for the two main LSPs, **lsp1-r6-to-r0** and **lsp2-r1-to-r5**. Interface **fe-0/1/0.0** is the connecting interface between **R1** and **R2**, and both LSPs are configured with a strict path through **fe-0/1/0.0**. For more detailed information about what is happening on interface **fe-0/1/0.0**, issue the **show rsvp interface extensive** command.

Sample Output

```

user@R1> show rsvp interface extensive
RSVP interface: 3 active

```



```

fe-0/1/0.0 Index 67, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 10.0.12.13
ActiveResv 2, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 100Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Protection: On, Bypass: 2, LSP: 2, Protected LSP: 2, Unprotected LSP: 0
  2 Jul 14 14:49:40 New bypass Bypass->10.0.12.14
  1 Jul 14 14:49:34 New bypass Bypass->10.0.12.14->10.0.24.2
Bypass: Bypass->10.0.12.14, State: Up, Type: LP, LSP: 0, Backup: 0
  3 Jul 14 14:49:42 Record Route: 10.0.17.14 10.0.27.1
  2 Jul 14 14:49:42 Up
  1 Jul 14 14:49:42 CSPF: computation result accepted
Bypass: Bypass->10.0.12.14->10.0.24.2, State: Up, Type: NP, LSP: 2, Backup: 0
  4 Jul 14 14:50:04 Record Route: 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
  3 Jul 14 14:50:04 Up
  2 Jul 14 14:50:04 CSPF: computation result accepted
  1 Jul 14 14:49:34 CSPF failed: no route toward 10.0.24.2
[...Output truncated...]

```

Meaning

Sample output from **R1** for the **show rsvp interface extensive** command shows more detailed information about the activity on all RSVP interfaces (3). However, only output for **fe-0/1/0.0** is shown. Protection is enabled (**Protection: On**), with two bypass paths (**Bypass: 2**) protecting two LSPs (**Protected LSP: 2**). All LSPs are protected, as indicated by the **Unprotected LSP: 0** field. The first bypass **Bypass->10.0.12.14** is a link protection bypass path (**Type: LP**), protecting the link between **R1** and **R2 fe-0/1/0.0**. The second bypass path **10.0.12.14->10.0.24.2** is a node-link protected LSP, avoiding **R2** in case of node failure.

Sample Output

```

user@R1> show rsvp session detail
Ingress RSVP: 2 sessions

192.168.4.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14->10.0.24.2
  Suggested label received: -, Suggested label sent: -

```


Recovery label received: -, Recovery label sent: 102000
 Resv style: 1 SE, Label in: -, Label out: 102000
 Time left: -, Since: Tue Jul 11 14:30:53 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 60120 protocol 0

Type: Bypass LSP

Number of data route tunnel through: 2

Number of RSVP session tunnel through: 0
 PATH rcvfrom: localclient
 Adspec: sent MTU 1500
 Path MTU: received 1500
 PATH sentto: 10.0.17.14 (fe-0/1/1.0) 336 pkts
 RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 310 pkts
Explct route: 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
 Record route: <self> 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1

192.168.5.1

From: 192.168.1.1, **LSPstate: Up**, ActiveRoute: 0
LSPname: lsp2-r1-to-r5, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 101872
 Resv style: 1 SE, Label in: -, Label out: 101872
 Time left: -, Since: Tue Jul 11 14:28:28 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 2 receiver 60118 protocol 0

Node/Link protection desired

Type: Node/Link protected LSP

PATH rcvfrom: localclient
 Adspec: sent MTU 1500
 Path MTU: received 1500
 PATH sentto: 10.0.12.14 (fe-0/1/0.0) 344 pkts
 RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 349 pkts
Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
 Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
 Total 2 displayed, Up 2, Down 0

Egress RSVP: 1 sessions

192.168.1.1

From: 192.168.5.1, **LSPstate: Up**, ActiveRoute: 0
LSPname: r5-to-r1, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: -


```

Resv style: 1 FF, Label in: 3, Label out: -
Time left: 147, Since: Tue Jul 11 14:28:36 2006
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 29228 protocol 0
PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 348 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>
Total 1 displayed, Up 1, Down 0

```

Transit RSVP: 2 sessions

192.168.0.1

```

From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
LSPname: lsp1-r6-to-r0, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 101952
Resv style: 1 SE, Label in: 100464, Label out: 101952
Time left: 134, Since: Tue Jul 11 14:31:38 2006
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 11131 protocol 0

```

Node/Link protection desired

Type: Node/Link protected LSP

```

PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 488 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.12.14 (fe-0/1/0.0) 339 pkts
RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 343 pkts

```

Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

```

Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

```

192.168.6.1

```

From: 192.168.0.1, LSPstate: Up, ActiveRoute: 0
LSPname: r0-to-t6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100448, Label out: 3
Time left: 158, Since: Tue Jul 11 14:31:36 2006
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 23481 protocol 0
PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 344 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.16.2 (so-0/0/3.0) 337 pkts

```



```

RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 310 pkts
Explot route: 10.0.16.2
Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning

Sample output from **R1** shows detailed information about the RSVP sessions active on **R1**. All sessions are up, with two ingress sessions, one egress session, and two transit sessions.

Within the ingress section, the first session is a bypass path, as indicated by the **Type: Bypass LSP** field; and the second session is a protected LSP (**lsp2-r1-to-r5**) originating on **R1**, as indicated by the **Type: Node/Link protected LSP** field.

Conclusion

Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection and node-link protection are facility-based methods used to reduce the amount of time needed to reroute LSP traffic. These protection methods are often compared to fast reroute—the other Junos OS LSP protection method.

While fast reroute protects LSPs on a one-to-one basis, link protection and node-link protection protect multiple LSPs by using a single, logical bypass LSP. Link protection provides robust backup support for a link, node-link protection bypasses a node or a link, and both types of protection are designed to interoperate with other vendor equipment. Such functionality makes link protection and node-link protection excellent choices for scalability, redundancy, and performance in MPLS-enabled networks.

Related Information

For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- *Junos User Guide*
- *Junos MPLS Applications Configuration Guide*
- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering*. White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

Verify That Link Protection Is Up

Purpose

When you verify link protection, you must check that the bypass LSP is up. You can also check the number of LSPs protected by the bypass. In the network shown in [Figure 147 on page 1906](#), a bypass path should

be up to protect the link between **R1** and **R2**, or next-hop **10.0.12.14**, and the two LSPs traversing the link, **lsp2-r1-to-r5** and **lsp1-r6-to-r0**.

Action

To verify link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router:

```
user@host> show mpls lsp extensive
user@host> show rsvp session detail
user@host> show rsvp interface
```

Sample Output

```
user@R1> show mpls lsp extensive | no-more
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: lsp2-r1-to-r5
  ActivePath:   via-r2 (primary)
Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary   via-r2           State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
      10.0.12.14(Label=101264) 10.0.24.2(Label=100736) 10.0.45.2(Label=3)
  6 Jun 16 14:06:33   Link-protection Up
  5 Jun 16 14:05:39 Selected as active path
  4 Jun 16 14:05:39 Record Route:  10.0.12.14(Label=101264)
10.0.24.2(Label=100736) 10.0.45.2(Label=3)
  3 Jun 16 14:05:39 Up
  2 Jun 16 14:05:39 Originate Call
  1 Jun 16 14:05:39 CSPF: computation result accepted
  Created: Fri Jun 16 14:05:38 2006
Total 1 displayed, Up 1, Down 0

[...Output truncated...]
```



```
Transit LSP: 2 sessions
```

```
192.168.0.1
```

```
From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
```

```
LSPname: lsp1-r6-to-r0, LSPpath: Primary
```

```
Suggested label received: -, Suggested label sent: -
```

```
Recovery label received: -, Recovery label sent: 101296
```

```
Resv style: 1 SE, Label in: 100192, Label out: 101296
```

```
Time left: 116, Since: Mon Jun 19 10:26:32 2006
```

```
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
```

```
Port number: sender 1 receiver 58739 protocol 0
```

```
Link protection desired
```

```
Type: Link protected LSP, using Bypass->10.0.12.14
```

```
1 Jun 19 10:26:32 Link protection up, using Bypass->10.0.12.14
```

```
PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 579 pkts
```

```
Adspec: received MTU 1500 sent MTU 1500
```

```
PATH sentto: 10.0.12.14 (fe-0/1/0.0) 474 pkts
```

```
RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 501 pkts
```

```
Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
```

```
Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
```

```
[...Output truncated...]
```

Meaning

The sample output from ingress router **R1** shows that **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have link protection up, and both LSPs are using the bypass path, **10.0.12.14**. However, the **show mpls lsp** command does not list the bypass path. For information about the bypass path, use the **show rsvp session** command.

Sample Output

```
user@R1> show rsvp session detail
```

```
Ingress RSVP: 2 sessions
```

```
192.168.2.1
```

```
From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
```

```
LSPname: Bypass->10.0.12.14
```

```
Suggested label received: -, Suggested label sent: -
```

```
Recovery label received: -, Recovery label sent: 101456
```

```
Resv style: 1 SE, Label in: -, Label out: 101456
```

```
Time left: -, Since: Fri May 26 18:38:09 2006
```

```
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
```

```
Port number: sender 1 receiver 18709 protocol 0
```

```
Type: Bypass LSP
```


Number of data route tunnel through: 2

Number of RSVP session tunnel through: 0

PATH rcvfrom: localclient

Adspec: sent MTU 1500

Path MTU: received 1500

PATH sentto: 10.0.17.14 (fe-0/1/1.0) 51939 pkts

RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 55095 pkts

Explct route: 10.0.17.14 10.0.27.1

Record route: <self> 10.0.17.14 10.0.27.1

192.168.5.1

From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0

LSPname: lsp2-r1-to-r5, LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: 101264

Resv style: 1 SE, Label in: -, Label out: 101264

Time left: -, Since: Fri Jun 16 14:05:39 2006

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 1 receiver 18724 protocol 0

Link protection desired**Type: Link protected LSP**

PATH rcvfrom: localclient

Adspec: sent MTU 1500

Path MTU: received 1500

PATH sentto: 10.0.12.14 (fe-0/1/0.0) 8477 pkts

RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 8992 pkts

Explct route: 10.0.12.14 10.0.24.2 10.0.45.2

Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2

Total 2 displayed, Up 2, Down 0

Egress RSVP: 1 sessions

192.168.1.1

From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0

LSPname: r5-to-r1, LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: -

Resv style: 1 FF, Label in: 3, Label out: -

Time left: 159, Since: Mon May 22 22:08:16 2006

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 1 receiver 64449 protocol 0

PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 63145 pkts

Adspec: received MTU 1500

PATH sentto: localclient


```

    RESV rcvfrom: localclient
    Record route: 10.0.59.1 10.0.79.2 10.0.17.14 <self>
Total 1 displayed, Up 1, Down 0
Transit RSVP: 2 sessions
192.168.0.1
    From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
    LSPname: lsp1-r6-to-r0, LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: 101296
    Resv style: 1 SE, Label in: 100192, Label out: 101296
    Time left: 129, Since: Mon Jun 19 10:26:32 2006
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 1 receiver 58739 protocol 0
    Link protection desired
    Type: Link protected LSP
    PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 3128 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.12.14 (fe-0/1/0.0) 2533 pkts
    RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 2685 pkts
    Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
    Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
    From: 192.168.0.1, LSPstate: Up, ActiveRoute: 0
    LSPname: r0-to-r6, LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: 3
    Resv style: 1 FF, Label in: 100128, Label out: 3
    Time left: 143, Since: Thu May 25 12:30:26 2006
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 1 receiver 4111 protocol 0
    PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 57716 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.16.2 (so-0/0/3.0) 54524 pkts
    RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 50534 pkts
    Explct route: 10.0.16.2
    Record route: 10.0.50.2 10.0.59.1 10.0.79.2 10.0.17.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning

The sample output from ingress router **R1** shows the ingress, egress, and transit LSPs for **R1**. Some information is similar to that found in the **show mpls lsp** command. However, because link protection is

an RSVP feature, information about bypass paths is provided. The bypass path appears as a separate RSVP ingress session for the protected interface, as indicated by the **Type** field.

The bypass path name is automatically generated. By default, the name appears as **Bypass > interface-address**, where the interface address is the next downstream router's interface (**10.0.12.14**). The explicit route **10.0.17.14 10.0.27.1** for the session shows **R7** as the transit node and **R2** as the egress node.

Within the ingress RSVP section of the output, the LSP originating at **R1 (lsp2-r1-to-r5)** is shown requesting link protection. Since a bypass path is in place to protect the downstream link, **lsp2-r1-to-r5** is associated with the bypass, as indicated by the **Link protected LSP** field.

The egress section of the output shows the return LSP **r5-to-r1**, which is not protected.

The transit section of the output shows link protection requested by **lsp1-r6-to-r0**. Since a bypass path is in place to protect the downstream link, **lsp1-r6-to-r0** is associated with the bypass, as indicated by the **Link protected LSP** field. Also included in the transit section of the output is the return LSP **r0-to-r6**, which is not protected.

Sample Output

```
user@R1> show rsvp interface
RSVP interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
fe-0/1/0.0	Up	2	100%	100Mbps	100Mbps	0bps	35Mbps
fe-0/1/1.0	Up	1	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/2.0	Up	0	100%	100Mbps	100Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning

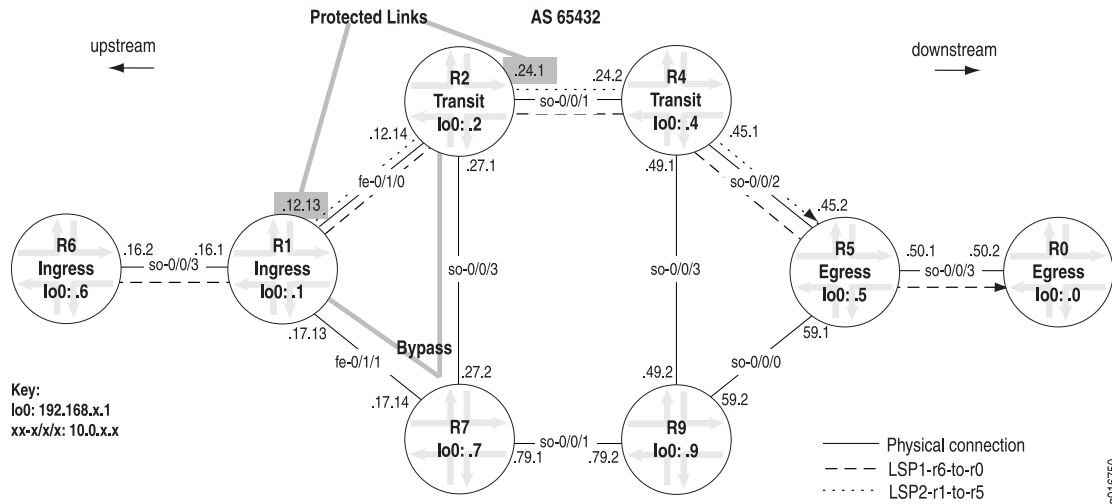
The sample output from ingress router **R1** shows the number of LSPs going through the interfaces configured on **R1**. The **Active resv** field shows the number of LSPs for each interface. For example, interface **fe-0/1/0.0** between **R1** and **R2** has two active reservations, **lsp1-r6-to-r0** and **lsp2-r1-to-r5**; interface **fe-0/1/1.0** between **R1** and **R7** has one, the bypass (**10.0.12.14**); interface **fe-0/1/2.0** between **R6** and **R1** has no LSP reservations; and interface **so-0/0/3.0** between **R6** and **R1** has one LSP reservation, **lsp1-r6-to-r0**.

Many-to-One Link Protection (Facility Backup) Overview

Many-to-one (facility backup) is based on interface rather than on LSP. While fast reroute protects interfaces or nodes along the entire path of a LSP, many-to-one protection can be applied on interfaces as needed, as shown in [Figure 147 on page 1906](#). In [Figure 147 on page 1906](#), a bypass path is set up around the link to

be protected (10.0.12.14) using an alternate interface to forward traffic. The bypass path is shared by all protected LSPs traversing the failed link (many LSPs protected by one bypass path).

Figure 147: Many-to-One or Link Protection



In Figure 147 on page 1906, two LSPs (lsp1-r6-to-r0 and lsp2-r1-to-r5) are protected by one preestablished bypass path from R1 to R2 through R7. Both LSPs have strict paths configured that go through interface fe-0/1/0. On R1, the interface 10.0.12.13 has link protection configured that protects the next hop 10.0.12.14.

Link protection (many-to-one or facility backup) allows a router immediately upstream from a link failure to use an alternate interface to forward traffic to its downstream neighbor. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link. A single bypass path can safeguard a set of protected LSPs. When an outage occurs, the router immediately upstream from the link outage switches protected traffic to the bypass link, then signals the link failure to the ingress router.

Like fast reroute, link protection provides local repair and restores connectivity faster than the ingress router switching traffic to a standby secondary path. However, unlike fast reroute, link protection does not provide protection against the failure of the downstream neighbor.

Link protection is appropriate in the following situations:

- The number of LSPs to be protected is large.
- Satisfying path selection criteria (priority, bandwidth, and link coloring) for bypass paths is less critical.
- Control at the granularity of individual LSPs is not required.

Verify One-to-One Backup

Purpose

You can verify that one-to-one backup is established by examining the ingress router and the other routers in the network.

Action

To verify one-to-one backup, enter the following Junos OS CLI operational mode commands:

```
user@host> show mpls lsp ingress extensive
user@host> show RSVP session
```

Sample Output

The following sample output is from the ingress router **R1** :

```
user@R1> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
    FastReroute desired
    LoadBalance: Random
    Encoding type: Packet, Switching type: Packet, GPID: IPv4
    *Primary   via-r2           State: Up
      SmartOptimizeTimer: 180
      Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          10.0.12.14(flag=9) 10.0.24.2(flag=1) 10.0.45.2
    8 May 11 14:51:46 Fast-reroute Detour Up
    7 May 11 14:50:55 Record Route: 10.0.12.14(flag=9) 10.0.24.2(flag=1) 10.0.45.2
6 May 11 14:50:55 Record Route: 10.0.12.14(flag=9) 10.0.24.2 10.0.45.2
5 May 11 14:50:52 Selected as active path
4 May 11 14:50:52 Record Route: 10.0.12.14 10.0.24.2 10.0.45.2
3 May 11 14:50:52 Up
2 May 11 14:50:52 Originate Call
1 May 11 14:50:52 CSPF: computation result accepted
```



```
Created: Thu May 11 14:50:52 2006
Total 1 displayed, Up 1, Down 0
```

Meaning

The sample output from **R1** shows that the **FastReroute desired** object was included in the Path messages for the LSP, allowing **R1** to select the active path for the LSP and establish a detour path to avoid **R2**.

In line 8, **Fast-reroute Detour Up** shows that the detour is operational. Lines 6 and 7 indicate that transit routers **R2** and **R4** have established their detour paths.

R2, 10.0.12.14, includes (**flag=9**), indicating that node protection is available for the downstream node and link. **R4, 10.0.24.2**, includes (**flag=1**), indicating that link protection is available for the next downstream link. In this case, **R4** can protect only the downstream link because the node is the egress router **R5**, which cannot be protected. For more information about flags, see the *Junos User Guide*.

The output for the **show mpls lsp extensive** command does not show the actual path of the detour. To see the actual links used by the detour paths, you must use the **show rsvp session ingress detail** command.

Sample Output

The following sample output is from the ingress router **R1** in the network shown in [Figure 163 on page 2163](#).

```
user@R1> show rsvp session ingress detail
Ingress RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1,  LSPstate: Up,  ActiveRoute: 0
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100848
  Resv style: 1 FF, Label in: -, Label out: 100848
  Time left:      -, Since: Thu May 11 14:17:15 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 9228 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 35 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 25 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
  Detour is Up
```



```

Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Detour adspec: sent MTU 1500
Path MTU: received 1500
Detour PATH sentto: 10.0.17.14 (fe-0/1/1.0) 23 pkts
Detour RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 20 pkts
Detour Explct route: 10.0.17.14 10.0.79.2 10.0.59.1
Detour Record route: <self> 10.0.17.14 10.0.79.2 10.0.59.1
Detour Label out: 100848
Total 1 displayed, Up 1, Down 0

```

Meaning

The sample output from **R1** shows the RSVP session of the main LSP. The detour path is established, **Detour is Up**. The physical path of the detour is displayed in **Detour Explct route**. The detour path uses **R7** and **R9** as transit routers to reach **R5**, the egress router.

Sample Output

The following sample output is from the first transit router **R2** in the network shown in [Figure 163 on page 2163](#):

```

user@R2> show rsvp session transit detail
Transit RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100448
  Resv style: 1 FF, Label in: 100720, Label out: 100448
  Time left: 126, Since: Wed May 10 16:12:21 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
FastReroute desired
  PATH rcvfrom: 10.0.12.13 (fe-0/1/0.0) 173 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.24.2 (so-0/0/1.0) 171 pkts
  RESV rcvfrom: 10.0.24.2 (so-0/0/1.0) 169 pkts
  Explct route: 10.0.24.2 10.0.45.2
  Record route: 10.0.12.13 <self> 10.0.24.2 10.0.45.2
Detour is Up
  Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Detour adspec: received MTU 1500 sent MTU 1500

```



```

Path MTU: received 1500
Detour PATH sentto: 10.0.27.2 (so-0/0/3.0) 169 pkts
Detour RESV rcvfrom: 10.0.27.2 (so-0/0/3.0) 167 pkts
Detour Explt route: 10.0.27.2 10.0.79.2 10.0.59.1
Detour Record route: 10.0.12.13 <self> 10.0.27.2 10.0.79.2 10.0.59.1
Detour Label out: 100736
Total 1 displayed, Up 1, Down 0

```

Meaning

The sample output from **R2** shows the detour is established (**Detour is Up**) and avoids **R4**, and the link connecting **R4** and **R5 (10.0.45.2)**. The detour path is through **R7 (10.0.27.2)** and **R9 (10.0.79.2)** to **R5 (10.0.59.1)**, which is different from the explicit route for the detour from **R1**. **R1** has the detour passing through the **10.0.17.14** link on **R7**, while **R1** is using the **10.0.27.2** link. Both detours merge at **R9** through the **10.0.79.2** link to **R5 (10.0.59.1)**.

Sample Output

The following sample output is from the second transit router **R4** in the network shown in [Figure 163 on page 2163](#):

```

user@R4> show rsvp session transit detail
Transit RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
    LSPname: r1-to-r5, LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: 3
    Resv style: 1 FF, Label in: 100448, Label out: 3
    Time left: 155, Since: Wed May 10 16:15:38 2006
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 5 receiver 9216 protocol 0
    FastReroute desired
    PATH rcvfrom: 10.0.24.1 (so-0/0/1.0) 178 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.45.2 (so-0/0/2.0) 178 pkts
    RESV rcvfrom: 10.0.45.2 (so-0/0/2.0) 175 pkts
    Explt route: 10.0.45.2
    Record route: 10.0.12.13 10.0.24.1 <self> 10.0.45.2
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: received MTU 1500 sent MTU 1500

```



```

Path MTU: received 1500
Detour PATH sentto: 10.0.49.2 (so-0/0/3.0) 176 pkts
Detour RESV rcvfrom: 10.0.49.2 (so-0/0/3.0) 175 pkts
Detour Explt route: 10.0.49.2 10.0.59.1
Detour Record route: 10.0.12.13 10.0.24.1 <self> 10.0.49.2 10.0.59.1
Detour Label out: 100352
Total 1 displayed, Up 1, Down 0

```

Meaning

The sample output from **R4** shows the detour is established (**Detour is Up**) and avoids the link connecting **R4** and **R5 (10.0.45.2)**. The detour path is through **R9 (10.0.49.2)** to **R5 (10.0.59.1)**. Some of the information is similar to that found in the output for **R1** and **R2**. However, the explicit route for the detour is different, going through the link connecting **R4** and **R9 (so-0/0/3 or 10.0.49.2)**.

Sample Output

The following sample output is from **R7**, which is used in the detour path in the network shown in [Figure 163 on page 2163](#):

```

user@R7> show rsvp session transit detail
Transit RSVP: 1 sessions, 1 detours

192.168.5.1
From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
LSPname: r1-to-r5, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 100368
Resv style: 1 FF, Label in: 100736, Label out: 100368
Time left: 135, Since: Wed May 10 16:14:42 2006
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 5 receiver 9216 protocol 0
Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0
PATH rcvfrom: 10.0.27.1 (so-0/0/3.0) 179 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.79.2 (so-0/0/1.0) 177 pkts
RESV rcvfrom: 10.0.79.2 (so-0/0/1.0) 179 pkts
Explt route: 10.0.79.2 10.0.59.1
Record route: 10.0.12.13 10.0.27.1 <self> 10.0.79.2 10.0.59.1
Label in: 100736, Label out: 100368

```


Detour branch from 10.0.17.13, to skip 192.168.2.1, Up

```

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0
PATH rcvfrom: 10.0.17.13 (fe-0/1/1.0) 179 pkts
Adspec: received MTU 1500
PATH sentto: 10.0.79.2 (so-0/0/1.0) 0 pkts
RESV rcvfrom: 10.0.79.2 (so-0/0/1.0) 0 pkts
Explot route: 10.0.79.2 10.0.59.1
Record route: 10.0.17.13 <self> 10.0.79.2 10.0.59.1

```

Label in: 100752, Label out: 100368

Total 1 displayed, Up 1, Down 0

Meaning

The sample output from **R7** shows the same information as for a regular transit router used in the primary path of the LSP: the ingress address (**192.168.1.1**), the egress address (**192.168.5.1**), and the name of the LSP (**r1-to-r5**). Two detour paths are displayed; the first to avoid **R4** (**192.168.4.1**) and the second to avoid **R2** (**192.168.2.1**). Because **R7** is used as a transit router by **R2** and **R4**, **R7** can merge the detour paths together as indicated by the identical **Label out** value (**100368**) for both detour paths. Whether **R7** receives traffic from **R4** with a label value of **100736** or from **R2** with a label value of **100752**, **R7** forwards the packet to **R5** with a label value of **100368**.

Sample Output

The following sample output is from **R9**, which is a router used in the detour path in the network shown in [Figure 163 on page 2163](#):

```
user@R9> show rsvp session transit detail
```

```
Transit RSVP: 1 sessions, 1 detours
```

```
192.168.5.1
```

```

From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
LSPname: r1-to-r5, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100352, Label out: 3
Time left: 141, Since: Wed May 10 16:16:40 2006
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 5 receiver 9216 protocol 0

```

Detour branch from 10.0.49.1, to skip 192.168.5.1, Up

```

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500

```



```

Path MTU: received 0
PATH rcvfrom: 10.0.49.1 (so-0/0/3.0) 183 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.59.1 (so-0/0/0.0) 182 pkts
RESV rcvfrom: 10.0.59.1 (so-0/0/0.0) 183 pkts
Explct route: 10.0.59.1
  Record route: 10.0.12.13 10.0.24.1 10.0.49.1 <self> 10.0.59.1
Label in: 100352, Label out: 3
Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0
Detour branch from 10.0.17.13, to skip 192.168.2.1, Up
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0
PATH rcvfrom: 10.0.79.1 (so-0/0/1.0) 181 pkts
Adspec: received MTU 1500
PATH sentto: 10.0.59.1 (so-0/0/0.0) 0 pkts
RESV rcvfrom: 10.0.59.1 (so-0/0/0.0) 0 pkts
Explct route: 10.0.59.1
Record route: 10.0.12.13 10.0.27.1 10.0.79.1 <self> 10.0.59.1
Label in: 100368, Label out: 3
Total 1 displayed, Up 1, Down 0

```

Meaning

The sample output from **R9** shows that **R9** is the penultimate router for the detour path, the explicit route includes only the egress link address (**10.0.59.1**), and the **Label out** value (**3**) indicates that **R9** has performed penultimate-hop label popping. Also, the detour branch from **10.0.27.1** does not include path information because **R7** has merged the detour paths from **R2** and **R4**. Notice that the **Label out** value in the detour branch from **10.0.17.13** is **100368**, the same value as the **Label out** value on **R7**.

Sample Output

The following sample output is from the egress router **R5** in the network shown in [Figure 163 on page 2163](#):

```

user@R5> show rsvp session egress detail
Egress RSVP: 1 sessions, 1 detours

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-to-r5, LSPpath: Primary

```



```

Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 119, Since: Thu May 11 14:44:31 2006
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 9230 protocol 0

```

FastReroute desired

```

PATH rcvfrom: 10.0.45.1 (so-0/0/2.0) 258 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient

```

Record route: 10.0.12.13 10.0.24.1 10.0.45.1 <self>

Detour branch from 10.0.49.1, to skip 192.168.5.1, Up

```

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0

```

Detour branch from 10.0.27.1, to skip 192.168.4.1, Up

```

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0

```

Detour branch from 10.0.17.13, to skip 192.168.2.1, Up

```

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Adspec: received MTU 1500
Path MTU: received 0

```

```

PATH rcvfrom: 10.0.59.2 (so-0/0/0.0) 254 pkts

```

```

Adspec: received MTU 1500

```

```

PATH sentto: localclient

```

```

RESV rcvfrom: localclient

```

```

Record route: 10.0.12.13 10.0.24.1 10.0.49.1 10.0.59.2 <self>

```

Label in: 3, Label out: -

```

Total 1 displayed, Up 1, Down 0

```

Meaning

The sample output from **R5** shows the main LSP in the **Record route** field and the detours through the network.

Verify That the Primary Path Is Operational

Purpose

Primary paths must always be used in the network if they are available, therefore an LSP always moves back to the primary path after a failure, unless the configuration is adjusted. For more information on

adjusting the configuration to prevent a failed primary path from reestablishing, see [“Preventing Use of a Path That Previously Failed”](#) on page 150.

Action

To verify that the primary path is operational, enter the following Junos OS command-line interface (CLI) operational mode commands:

```
user@host> show mpls lsp extensive ingress
user@host> show rsvp interface
```

Sample Output 1

```
user@R1> show mpls lsp extensive ingress
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
    ActivePath: via-r2 (primary)
    LoadBalance: Random
    Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary via-r2           State: Up
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 11)
10.0.12.14 S 10.0.24.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          10.0.12.14 10.0.24.2
5 Apr 29 14:40:43 Selected as active path
4 Apr 29 14:40:43 Record Route: 10.0.12.14 10.0.24.2
3 Apr 29 14:40:43 Up
2 Apr 29 14:40:43 Originate Call
1 Apr 29 14:40:43 CSPF: computation result accepted
Standby via-r7           State: Dn
  SmartOptimizeTimer: 180
    No computed ERO.
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0
```


Sample Output 2

```
user@R1> show rsvp interface
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
fe-0/1/0.0	Up	2	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/1.0	Up	1	100%	100Mbps	100Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning

Sample output 1 shows that the LSP is operational and is using the primary path (**via-r2**) with **R2 (10.0.12.14)** and **R4 (10.0.24.2)** as transit routers. The priority values are the same for setup and hold, **6 6**. Priority 0 is the highest (best) priority and 7 is the lowest (worst) priority. The Junos OS default for setup and hold priority is 7:0. Unless some LSPs are more important than others, preserving the default is a good practice. Configuring a setup priority that is better than the hold priority is not allowed, resulting in a failed commit in order to avoid preemption loops.

Verify That the Secondary Path Is Established

Purpose

When the secondary path is configured with the **standby** statement, the secondary path should be *up* but *not active*; it will become active if the primary path fails. A secondary path configured without the **standby** statement will not come up unless the primary path fails. To test that the secondary path is correctly configured and would come up if the primary path were to fail, you must deactivate a link or node critical to the primary path, then issue the **show mpls lsp *lsp-path-name* extensive** command.

Action

To verify that the secondary path is established, enter the following Junos OS CLI operational mode command:

Sample Output

```
user@R1> show mpls lsp extensive
```


Sample Output

The following sample output shows a correctly configured secondary path before and after it comes up. In the example, interface **fe-0/1/0** on **R2** is deactivated, which brings down the primary path **via-r2**. The ingress router **R1** switches traffic to the secondary path **via-r7**.

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-r2                State: Up
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.0.12.14 10.0.24.2 10.0.45.2
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route: 10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
  Secondary via-r7                State: Dn
    SmartOptimizeTimer: 180
    No computed ERO.
    Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0

[edit interfaces]
user@R2# deactivate fe-0/1/0

[edit interfaces]
user@R2# show
inactive: fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.12.14/30;

```



```

    }
    family iso;
    family mpls;
  }
}

```

user@R1> **show mpls lsp name r1-to-r4 extensive**

Ingress LSP: 1 sessions

192.168.4.1

From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r4

ActivePath: via-r7 (secondary)

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

Primary via-r2 State: Dn

Priorities: 6 6

Bandwidth: 35Mbps

SmartOptimizeTimer: 180

Will be enqueued for recomputation in 14 second(s).

10 Apr 29 14:52:33 **CSPF failed: no route toward 10.0.12.1** 4[21 times]

9 Apr 29 14:42:48 Clear Call

8 Apr 29 14:42:48 Deselected as active

7 Apr 29 14:42:48 Session preempted

6 Apr 29 14:42:48 Down

5 Apr 29 14:40:43 Selected as active path

4 Apr 29 14:40:43 Record Route: 10.0.12.14 10.0.24.2

3 Apr 29 14:40:43 Up

2 Apr 29 14:40:43 Originate Call

1 Apr 29 14:40:43 CSPF: computation result accepted

***Standby via-r7 State: Up**

SmartOptimizeTimer: 180

Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 11)

10.0.17.14 S 10.0.47.1 S

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.0.17.14 10.0.47.1

5 Apr 29 14:42:48 Selected as active path

4 Apr 29 14:41:12 Record Route: 10.0.17.14 10.0.47.1

3 Apr 29 14:41:12 Up

2 Apr 29 14:41:12 Originate Call

1 Apr 29 14:41:12 CSPF: computation result accepted

Created: Sat Apr 29 14:40:43 2006

Total 1 displayed, Up 1, Down 0

Meaning

The sample output from egress router **R1** shows a correctly configured standby secondary path in a down state because the primary path is still up. Upon deactivation of an interface (**interface fe-0/1/0** on **R2**) critical to the primary path, the primary path **via-r2** goes down and the standby secondary path **via-r7** comes up, allowing **R1** to switch traffic to the standby secondary path.

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To verify that the LSP is up, enter some or all of the following commands from the ingress router:

```
user@host> show mpls lsp
user@host> show mpls lsp extensive
user@host> show mpls lsp name name
user@host> show mpls lsp name name extensive
```

Sample Output 1

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6     10.0.0.1      Dn      0  -              R1-to-R6
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp
Ingress LSP: 0 sessions
```


Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

user@R6> **show mpls lsp**

Ingress LSP: 1 sessions

To	From	State	Rt	ActivePath	P	LSPname
10.0.0.1	10.0.0.6	Dn	0	-		R6-to-R1

Total 1 displayed, Up 0, **Down 1**

Egress LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Sample Output 2

user@R1> **show mpls lsp extensive**

Ingress LSP: 1 sessions

10.0.0.6

From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6

ActivePath: (none)

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

Primary State: Dn

Will be enqueued for recomputation in 22 second(s).

1 Nov 2 14:43:38 **CSPF failed: no route toward 10.0.0.6** [175 times]

Created: Tue Nov 2 13:18:39 2004

Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions


```

Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 13 second(s).
    1 Nov  2 14:38:12  CSPF failed: no route toward 10.0.0.1 [177 times]
  Created: Tue Nov  2 13:12:22 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 3

```

user@R1> show mpls lsp name R1-to-R6
Ingress LSP: 1 sessions

```

To	From	State	Rt	ActivePath	P	LSPname
10.0.0.6	10.0.0.1	Dn	0	-		R1-to-R6

```

Total 1 displayed, Up 0, Down 1

```



```
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 4

```
user@R1> show mpls lsp name R1-to-R6 extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 10 second(s).
    1 Nov  2 14:51:53 CSPF failed: no route toward 10.0.0.6[192 times]
  Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

Sample Output 1 shows a brief description of the state of the LSP for the ingress, transit, and egress routers. Output from ingress router **R1** and egress router **R6** shows that both LSPs are down, **R1-to-R6** and **R6-to-R1**. With the configured LSPs on **R1** and **R6**, we would expect egress LSP sessions on both **R1** and **R6**. In addition, transit router **R3** has no transit sessions.

Sample Output 2 shows all information about the LSPs, including all past state history and the reason why an LSP failed. Output from **R1** and **R6** indicates that there is no route to the destination because the Constrained Shortest Path First (CSPF) algorithm failed.

Sample Outputs 3 and 4 show examples of the output for the **show mpls lsp name** command with the **extensive** option. In this instance, the output is very similar to the **show mpls lsp** command because only

one LSP is configured in the example network in [Figure 146 on page 1873](#). However, in a large network with many LSPs configured, the results would be quite different between the two commands.

Verify the LSP Route on the Transit Router

Purpose

If the LSP is up, the LSP route should appear in the **mpls.0** routing table. MPLS maintains an MPLS path routing table (**mpls.0**), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP. If routes are not present in the output for the transit router, check the MPLS protocol configuration on the ingress and egress routers.

Action

To verify the LSP route on the transit router, enter the following command from the transit router:

```
user@host> show route table mpls.0
```

Sample Output 1

```
user@R3> show route table mpls.0
```

```
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 16w2d 21:52:40, metric 1
           Receive
1          *[MPLS/0] 16w2d 21:52:40, metric 1
           Receive
2          *[MPLS/0] 16w2d 21:52:40, metric 1
           Receive
```

Sample Output 2

```
user@R3> show route table mpls.0
```

```
mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```



```

0          *[MPLS/0] 16w2d 22:26:08, metric 1
           Receive
1          *[MPLS/0] 16w2d 22:26:08, metric 1
           Receive
2          *[MPLS/0] 16w2d 22:26:08, metric 1
           Receive
100864     *[RSVP/7] 00:07:23, metric 1
           > via so-0/0/2.0, label-switched-path R6-to-R1
100864(S=0) *[RSVP/7] 00:07:23, metric 1
           > via so-0/0/2.0, label-switched-path R6-to-R1
100880     *[RSVP/7] 00:07:01, metric 1
           > via so-0/0/3.0, label-switched-path R1-to-R6
100880(S=0) *[RSVP/7] 00:07:01, metric 1
           > via so-0/0/3.0, label-switched-path R1-to-R6

```

Meaning

Sample Output 1 from transit router **R3** shows three route entries in the form of MPLS label entries. These MPLS labels are reserved MPLS labels defined in RFC 3032, and are always present in the **mpls.0** routing table, regardless of the state of the LSP. The incoming labels assigned by RSVP to the upstream neighbor are missing from the output, indicating that the LSP is down. For more information on MPLS label entries, see [“Checklist for Verifying LSP Use” on page 2090](#).

In contrast, Sample Output 2 shows the MPLS labels and routes for a correctly configured LSP. The three reserved MPLS labels are present, and the four other entries represent the incoming labels assigned by RSVP to the upstream neighbor. These four entries represent two routes. There are two entries per route because the stack values in the MPLS header may be different. For each route, the second entry **100864 (S=0)** and **100880 (S=0)** indicates that the stack depth is not 1, and additional label values are included in the packet. In contrast, the first entry, **100864** and **100880** has an inferred S=1 value which indicates a stack depth of 1 and makes each label the last label in that particular packet. The dual entries indicate that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

Verify the LSP Route on the Ingress Router

Purpose

Check whether the LSP route is included in the active entries in the **inet.3** routing table for the specified address.

Action

To verify the LSP route, enter the following command from the ingress router:


```
user@host> show route destination
```

Sample Output 1

```
user@R1> show route 10.0.0.6
```

```
inet.0 : 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.6/32      *[IS-IS/18] 6d 01:41:37, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
                  to 10.1.13.2 via so-0/0/2.0
```

```
user@R6> show route 10.0.0.1
```

```
inet.0 : 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.1/32      *[IS-IS/18] 5d 01:01:38, metric 20
                  to 10.1.56.1 via so-0/0/0.0
                  > to 10.1.26.1 via so-0/0/2.0
                  to 10.1.36.1 via so-0/0/3.0
```

Sample Output 2

```
user@R1> show route 10.0.0.6
```

```
inet.0: 28 destinations, 28 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32      *[IS-IS/18] 6d 02:13:42, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
                  to 10.1.13.2 via so-0/0/2.0

inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```



```

10.0.0.6/32          *[RSVP/7] 00:08:07, metric 20
                    > via so-0/0/2.0, label-switched-path R1-to-R6

user@R6> show route 10.0.0.1

inet.0: 29 destinations, 29 routes (28 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.1/32          *[IS-IS/18] 5d 01:34:03, metric 20
                    to 10.1.56.1 via so-0/0/0.0
                    > to 10.1.26.1 via so-0/0/2.0
                    to 10.1.36.1 via so-0/0/3.0

inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.1/32          *[RSVP/7] 00:10:39, metric 20
                    > via so-0/0/3.0, label-switched-path R6-to-R1

```

Meaning

Sample Output 1 shows entries in the **inet.0** routing table only. The **inet.3** routing table is missing from the output because the LSP is not working. The **inet.0** routing table is used by interior gateway protocols (IGPs) and Border Gateway Protocol (BGP) to store routing information. In this case, the IGP is Intermediate System-to-Intermediate System (IS-IS). For more information on the **inet.0** routing table, see the *Junos MPLS Applications Configuration Guide*.

If the LSP was working, we would expect to see entries that include the LSP in the **inet.3** routing table. The **inet.3** routing table is used on ingress routers to route BGP packets to the destination egress router. BGP uses the **inet.3** routing table on the ingress router to help resolve next-hop addresses. BGP is configured in the example network shown in [Figure 146 on page 1873](#).

Sample Output 2 shows output you should receive when the LSP is up. The output shows both the **inet.0** and **inet.3** routing tables, indicating that LSPs **R1-to-R6** and **R6-to-R1** are available.

Verify MPLS Labels with the traceroute Command

Purpose

Display the route packets take to a BGP destination where the BGP next hop for that route is the LSP egress address. By default, BGP uses the **inet.0** and the **inet.3** routing tables to resolve the next-hop address. When the next-hop address of the BGP route is not the router ID of the egress router, traffic is

mapped to IGP routes, not to the LSP. Use the **tracert** command as a debugging tool to determine whether the LSP is being used to forward traffic.

Action

To verify MPLS labels, enter the following command from the ingress router:

```
user@host> traceroute hostname
```

Sample Output 1

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.12.2 (10.1.12.2)  0.627 ms  0.561 ms  0.520 ms
 2  10.1.26.2 (10.1.26.2)  0.570 ms !N  0.558 ms !N  4.879 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.630 ms  0.545 ms  0.488 ms
 2  10.1.12.1 (10.1.12.1)  0.551 ms !N  0.557 ms !N  0.526 ms !N
```

Sample Output 2

```
user@R1> traceroute 100.100.6.1
to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.866 ms  0.746 ms  0.724 ms
    MPLS Label=100912 CoS=0 TTL=1 S=1
 2  10.1.36.2 (10.1.36.2)  0.577 ms !N  0.597 ms !N  0.546 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.36.1 (10.1.36.1)  0.802 ms  0.716 ms  0.688 ms
    MPLS Label=100896 CoS=0 TTL=1 S=1
 2  10.1.13.1 (10.1.13.1)  0.570 ms !N  0.568 ms !N  0.546 ms !N
```

Meaning

Sample Output 1 shows that BGP traffic is not using the LSP, consequently MPLS labels do not appear in the output. Instead of using the LSP, BGP traffic is using the IGP (IS-IS, in the example network in [Figure 146 on page 1873](#)) to reach the BGP next-hop LSP egress address. The Junos OS default behavior uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Sample Output 2 is an example of output for a correctly configured LSP. The output shows MPLS labels, indicating that BGP traffic is using the LSP to reach the BGP next hop.

Verify MPLS Labels with the ping Command

Purpose

When you ping a specific LSP, you check that echo requests are sent over the LSP as MPLS packets.

Action

To verify MPLS labels, enter the following command from the ingress router to ping the egress router:

```
user@host> ping mpls rsvp lsp-name detail
```

For example:

```
user@R1> ping mpls rsvp R1-to-R6 detail
```

Sample Output 1

```
user@R1> ping mpls rsvp R1-to-R6 detail
LSP R1-to-R6 - LSP has no active path, exiting.

user@R6> ping mpls rsvp R6-to-R1 detail
LSP R6-to-R1 - LSP has no active path, exiting.
```

Sample Output 2

```
user@R1> traceroute 10.0.0.6
traceroute to 10.0.0.6 (10.0.0.6), 30 hops max, 40 byte packets
```



```

1  10.1.15.2 (10.1.15.2)  0.708 ms  0.613 ms  0.576 ms
2  10.0.0.6 (10.0.0.6)  0.763 ms  0.708 ms  0.700 ms

user@R1> ping mpls rsvp R1-to-R6 detail
Request for seq 1, to interface 69, label 100880
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 69, label 100880
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 69, label 100880
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 69, label 100880
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 69, label 100880
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

user@R6> ping mpls rsvp R6-to-R1 detail
Request for seq 1, to interface 70, label 100864
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 70, label 100864
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 70, label 100864
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 70, label 100864
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 70, label 100864
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

```

Meaning

Sample Output 1 shows that the LSP does not have an active path to forward echo requests, indicating that the LSP is down.

Sample Output 2 is an example of output you should receive when the LSP is up and forwarding packets.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, an interface is incorrectly configured at the `[edit protocols mpls]` hierarchy level on egress router **R6**.

Solution

To correct the error in this example, follow these steps:

1. Activate the interface in the MPLS protocol configuration on egress router **R6**:

```
user@R6> edit
user@R6# edit protocols mpls
[edit protocols mpls]
user@R6# show
user@R6# activate interface so-0/0/3.0
```

2. Verify and commit the configuration:

```
[edit protocols mpls]
user@R6# show
user@R6# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols mpls

[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured interface

[edit protocols mpls]
user@R6# activate interface so-0/0/3
```



```
[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
  to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
interface so-0/0/3.0; <<< Correctly configured interface

[edit protocols mpls]
user@R6# commit
commit complete
```

Meaning

The sample output shows that the incorrectly configured interface **so-0/0/3.0** on egress router **R6** is now activated at the **[edit protocols mpls]** hierarchy level. The LSP can now come up.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the BGP layer has been resolved.

Action

To verify the LSP again, enter the following command from the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1,  State: Up , ActiveRoute: 1 , LSPname: R1-to-R6
```



```

ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
        6 Nov  2 15:48:52 Selected as active path
        5 Nov  2 15:48:52 Record Route:  10.1.13.2 10.1.36.2
        4 Nov  2 15:48:52 Up
        3 Nov  2 15:48:52 Originate Call
        2 Nov  2 15:48:52 CSPF: computation result accepted
        1 Nov  2 15:48:22 CSPF failed: no route toward 10.0.0.6[308 times]
    Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6,  LSPstate: Up , ActiveRoute: 0
LSPname: R6-to-R1 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 159, Since: Tue Nov  2 15:48:30 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39106 protocol 0
PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions

```


Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1

From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
 LSPname: R6-to-R1, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100864, Label out: 3
 Time left: 123, Since: Tue Nov 2 15:35:41 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 39106 protocol 0
 PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.13.1 (so-0/0/2.0) 10 pkts
 RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
 Explct route: 10.1.13.1
 Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
 LSPname: R1-to-R6, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100880, Label out: 3
 Time left: 145, Since: Tue Nov 2 15:36:03 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 48015 protocol 0
 PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.36.2 (so-0/0/3.0) 10 pkts
 RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
 Explct route: 10.1.36.2
 Record route: 10.1.13.1 <self> 10.1.36.2

Total 2 displayed, Up 2, Down 0

user@R6> **show mpls lsp extensive**

Ingress LSP: 1 sessions

10.0.0.1

From: 10.0.0.6, **State: Up, ActiveRoute: 1, LSPname: R6-to-R1**
 ActivePath: (primary)
 LoadBalance: Random


```

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.36.1 10.1.13.1
        6 Nov  2 15:41:44 Selected as active path
        5 Nov  2 15:41:44 Record Route:  10.1.36.1 10.1.13.1
        4 Nov  2 15:41:44 Up
        3 Nov  2 15:41:44 Originate Call
        2 Nov  2 15:41:44 CSPF: computation result accepted
        1 Nov  2 15:41:14 CSPF failed: no route toward 10.0.0.1[306 times]
    Created: Tue Nov  2 13:12:21 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
    From: 10.0.0.1,  LSPstate: Up,  ActiveRoute: 0
    LSPname: R1-to-R6 , LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: -
    Resv style: 1 FF, Label in: 3, Label out: -
    Time left:  157, Since: Tue Nov  2 15:42:06 2004
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 1 receiver 48015 protocol 0
    PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
    Adspec: received MTU 1500
    PATH sentto: localclient
    RESV rcvfrom: localclient
    Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Checklist for Working with the Layered MPLS Troubleshooting Model

Problem

Description: This checklist provides a link to more detailed information about the layered Multiprotocol Label Switching network.

Solution

[Table 51 on page 1935](#) provides commands for working with the layered MPLS troubleshooting model.

Table 51: Checklist for Working with the Layered MPLS Troubleshooting Model

Tasks	Command or Action
“Understanding the Layered MPLS Troubleshooting Model” on page 1935	show mpls lsp show mpls lsp extensive show mpls lsp name <i>name</i> show mpls lsp name <i>name</i> extensive

Understanding the Layered MPLS Troubleshooting Model

Problem

Description: The layered MPLS troubleshooting model is a disciplined approach to investigating problems with an MPLS network. [Figure 148 on page 1936](#) illustrates the layers in the model, and the commands you can use to structure your investigation. Because of the complexity of the MPLS network, you can obtain much better results from your investigations if you progress through the layers and verify the functioning of each layer on the ingress, egress, and transit routers before moving on to the next layer.

Solution

[Figure 148 on page 1936](#) shows the layered MPLS troubleshooting model that you can use to troubleshoot problems with your MPLS network.

Figure 148: Layered MPLS Network Troubleshooting Model

BGP Layer	traceroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
<div>↩</div> IGP and IP Layers Functioning <div>↪</div>	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive <i>JUNOS Interfaces Network Operations Guide</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

g015528

As you move from one layer of the model to the next, you verify the correct functioning of a different component of the MPLS network and eliminate that layer as the source of the problem.

Physical Layer

When you investigate the physical layer, you check that the routers are connected, and the interfaces are up and configured correctly. To check the physical layer, enter the **show interfaces**, **show interfaces terse**, and **ping** commands. If there is a problem in the physical layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information on checking the physical layer, see [“Checklist for Verifying the Physical Layer” on page 1943](#).

Data Link Layer

When you investigate the data link layer, you check the encapsulation mode, for example, Point-to-Point Protocol (PPP) or Cisco High-Level Data Link Control (HDLC); PPP options, for example, header encapsulation; frame check sequence (FCS) size; and whether keepalive frames are enabled or disabled. To check the data link layer, enter the **show interfaces extensive** command. If there is a problem in the data link layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information on checking the data link layer, see [“Checking the Data Link Layer” on page 1961](#) and the *Junos Interfaces Operations Guide*.

IP Layer

When you investigate the IP layer, you verify that interfaces have correct IP addressing, and that the interior gateway protocol (IGP) neighbor adjacencies are established. To check the IP layer, enter the **show**

interfaces terse, **show ospf neighbor extensive**, and **show isis adjacency extensive** commands. If there is a problem in the IP layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.

IGP Layer

When you investigate the IGP layer, you verify that the the Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) protocols are configured and running correctly.

- If you have the OSPF protocol configured, you must check the IP layer first, and then the OSPF configuration. When you investigate the OSPF layer, you check that the protocol, interfaces, and traffic engineering are configured correctly. To check the OSPF layer, enter the **show configuration protocols ospf** and **show ospf interface** commands. If the problem exists in the OSPF layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information about checking the OSPF layer, see [“Verifying the OSPF Protocol” on page 2019](#).
- If you have the IS-IS protocol configured, because IS-IS and IP are independent of each other, it doesn't matter which one you check first. When you check the IS-IS configuration, you verify that IS-IS adjacencies are up, and the interfaces and IS-IS protocol are configured correctly. To check the IS-IS layer, enter the **show isis adjacency**, **show configuration protocols isis**, and **show isis interfaces** commands. If the problem exists in the IS-IS layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information about checking the IS-IS layer, see *Verifying the IS-IS Protocol*.

NOTE: The IS-IS protocol has traffic engineering enabled by default.

RSVP and MPLS Layers

After you have both the IP and IGP layers functioning and the problem is still not solved, you can begin to check the Resource Reservation Protocol (RSVP) and MPLS layers to determine if the problem is in one of these layers.

- When you investigate the RSVP layer, you are checking that dynamic RSVP signaling is occurring as expected, neighbors are connected, and interfaces are configured correctly for RSVP. To check the RSVP layer, enter the **show rsvp session**, **show rsvp neighbor**, and **show rsvp interface** commands. If there is a problem in the RSVP layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.
- When you investigate the MPLS layer, you are checking whether the LSP is up and functioning correctly. To check the MPLS layer, enter the **show mpls lsp**, **show mpls lsp extensive**, **show route table mpls.0**, **show route address**, **traceroute address**, and **ping mpls rsvp lsp-name detail** commands. If there is a problem in the MPLS layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.

BGP Layer

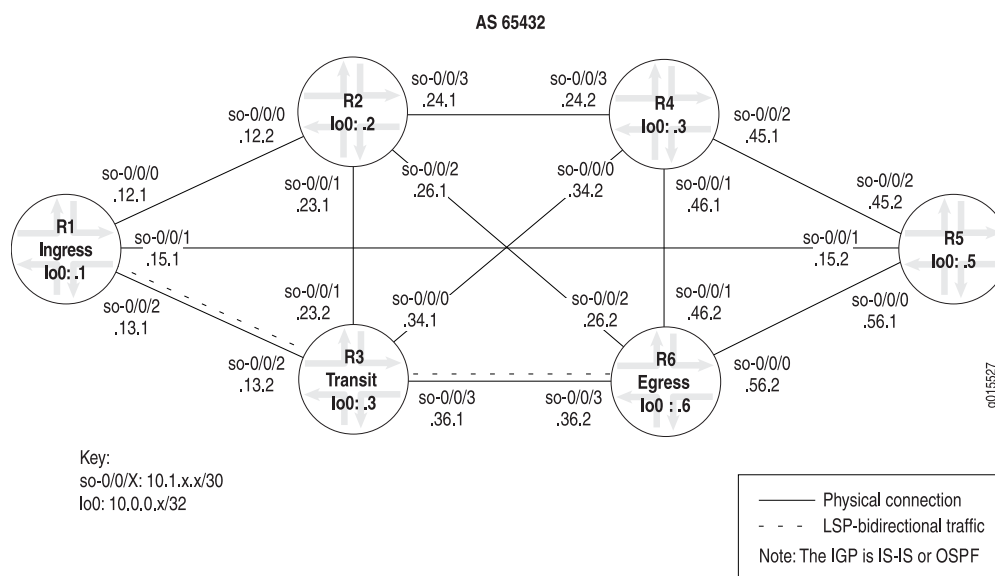
If the problem persists after you have checked the RSVP and MPLS layers, you must verify that the Border Gateway Protocol (BGP) is working correctly. There is no point in checking the BGP layer unless the LSP is established because BGP uses the MPLS LSP to forward traffic. When you check the BGP layer, you verify that the route is present and active, and more importantly, you ensure that the next hop is the LSP. To check the BGP layer, enter the **traceroute host-name**, **show bgp summary**, **show configuration protocols bgp**, **show route destination-prefix detail**, and **show route receive protocol bgp neighbor-address** commands. For more information on checking the BGP layer, see *Checking the BGP Layer*.

In reality, you could start at any level of the MPLS model to investigate a problem with your MPLS network. However, a disciplined approach, as the one described here, produces more consistent and reliable results.

Figure 149 on page 1938 illustrates the basic network topology used in the following topics that demonstrate how to troubleshoot an MPLS network:

- [Checklist for Verifying the Physical Layer on page 1943](#)
- [Checklist for Checking the Data Link Layer on page 1960](#)
- [Checklist for Verifying the IP and IGP Layers on page 1988](#)
- [Checklist for Checking the RSVP Layer on page 2055](#)
- [Checklist for Checking the MPLS Layer on page 1871](#)
- *Checklist for Checking the BGP Layer*

Figure 149: MPLS Basic Network Topology Example



The MPLS network consists of the following components:

- Router-only network with SONET interfaces
- MPLS protocol enabled on all routers, with interfaces selectively deactivated to illustrate a particular problem scenario
- All interfaces configured with MPLS
- A full-mesh IBGP topology, using AS 65432
- IS-IS or OSPF as the underlying IGP, using one level (IS-IS Level 2) or one area (OSPF area 0.0.0.0)
- A **send-statics** policy on routers R1 and R6, allowing a new route to be advertised into the network
- Two LSPs between routers R1 and R6, allowing for bidirectional traffic.

After you have configured an LSP, it is considered best practice to issue the **show mpls lsp** command to verify that the LSP is up, and to investigate further if you find an error message in the output. The error message can indicate a problem at any layer of the MPLS network.

The LSPs can be ingress, transit, or egress. Use the **show mpls lsp** command for quick verification of the LSP state, with the **extensive** option (**show mpls lsp extensive**) as a follow-up if the LSP is down. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To begin the investigation of an error in your MPLS network, from the ingress router, enter some or all of the following Junos OS command-line interface (CLI) operational mode commands:

```
user@host> show mpls lsp
user@host> show mpls lsp extensive
user@host> show mpls lsp name name
user@host> show mpls lsp name name extensive
```

Sample Output 1

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt  ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Up      1                *      R1-to-R6
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
10.0.0.1    10.0.0.6    Up      0 1 FF          3      - R6-to-R1
Total 1 displayed, Up 1, Down 0
```



```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
30 Dec 28 13:47:29 Selected as active path
29 Dec 28 13:47:29 Record Route: 10.1.13.2 10.1.36.2
28 Dec 28 13:47:29 Up
27 Dec 28 13:47:29 Originate Call
26 Dec 28 13:47:29 CSPF: computation result accepted
25 Dec 28 13:46:59 CSPF failed: no route toward 10.0.0.6
24 Dec 28 13:46:39 Deselected as active
23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6
22 Dec 28 13:46:39 Clear Call
21 Dec 28 13:46:39 ResvTear received
20 Dec 28 13:46:39 Down
19 Dec 28 13:46:39 10.1.13.2: Session preempted
18 Dec 28 13:42:07 Selected as active path
17 Dec 28 13:42:07 Record Route: 10.1.13.2 10.1.36.2
16 Dec 28 13:42:07 Up
15 Dec 28 13:42:07 Originate Call
14 Dec 28 13:42:07 CSPF: computation result accepted
13 Dec 28 13:41:37 CSPF failed: no route toward 10.0.0.6
12 Dec 28 13:41:16 Deselected as active
11 Dec 28 13:41:16 CSPF failed: no route toward 10.0.0.6
10 Dec 28 13:41:16 Clear Call
 9 Dec 28 13:41:16 ResvTear received
 8 Dec 28 13:41:16 Down
 7 Dec 28 13:41:16 10.1.13.2: Session preempted
```



```

6 Dec 13 11:50:15 Selected as active path
5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2
4 Dec 13 11:50:15 Up
3 Dec 13 11:50:15 Originate Call
2 Dec 13 11:50:15 CSPF: computation result accepted
1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]
---(more)---[abort]

```

Sample Output 3

```

user@R1> show mpls lsp name R1-to-R6
Ingress LSP: 1 sessions

```

To	From	State	Rt	ActivePath	P	LSPname
10.0.0.6	10.0.0.1	Up	1		*	R1-to-R6

```

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 4

```

user@R1> show mpls lsp name R1-to-R6 extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.13.2 10.1.36.2
30 Dec 28 13:47:29 Selected as active path
29 Dec 28 13:47:29 Record Route: 10.1.13.2 10.1.36.2
28 Dec 28 13:47:29 Up

```



```

27 Dec 28 13:47:29 Originate Call
26 Dec 28 13:47:29 CSPF: computation result accepted
25 Dec 28 13:46:59 CSPF failed: no route toward 10.0.0.6
24 Dec 28 13:46:39 Deselected as active
23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6
22 Dec 28 13:46:39 Clear Call
21 Dec 28 13:46:39 ResvTear received
20 Dec 28 13:46:39 Down
19 Dec 28 13:46:39 10.1.13.2: Session preempted
18 Dec 28 13:42:07 Selected as active path
17 Dec 28 13:42:07 Record Route: 10.1.13.2 10.1.36.2
16 Dec 28 13:42:07 Up
15 Dec 28 13:42:07 Originate Call
14 Dec 28 13:42:07 CSPF: computation result accepted
13 Dec 28 13:41:37 CSPF failed: no route toward 10.0.0.6
12 Dec 28 13:41:16 Deselected as active
11 Dec 28 13:41:16 CSPF failed: no route toward 10.0.0.6
10 Dec 28 13:41:16 Clear Call
9 Dec 28 13:41:16 ResvTear received
8 Dec 28 13:41:16 Down
7 Dec 28 13:41:16 10.1.13.2: Session preempted
6 Dec 13 11:50:15 Selected as active path
5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2
4 Dec 13 11:50:15 Up
3 Dec 13 11:50:15 Originate Call
2 Dec 13 11:50:15 CSPF: computation result accepted
1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]
Created: Mon Dec 13 11:47:19 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output from the ingress router **R1** shows that the label-switched path is traversing the network as intended, from **R1** through **R3** to **R6**, and another LSP in the reverse direction, from **R6** through **R3** to **R1**.

If your network has numerous LSPs, you might consider using the **show mpls lsp** command for quick verification of the LSP state. and the **show mpls lsp name *name* extensive** command to continue your investigation if you find that the LSP is down.

For more information about the status and statistics of the **show mpls lsp** command, see [“Checklist for Determining LSP Status” on page 2087](#). For more information on the availability and valid use of an LSP, see [“Checklist for Verifying LSP Use” on page 2090](#).

In all the following topics, the network topology is broken at different layers of the network so that you can investigate various MPLS network problems. The problems presented are not inclusive. Instead, the problems serve to illustrate one possible process of investigation into the different layers of the troubleshooting model.

RELATED DOCUMENTATION

[Verifying the Physical Layer | 1944](#)

[Checking the Data Link Layer | 1961](#)

[Verifying the IP and IGP Layers | 1990](#)

[Checking the RSVP Layer | 2056](#)

[Checking the MPLS Layer | 1872](#)

Checking the BGP Layer

Checklist for Verifying the Physical Layer

Problem

Description: This checklist provides the steps and commands for investigating a problem at the physical layer of a Multiprotocol Label Switching (MPLS) network. The checklist provides links to an overview of verifying the physical layer and more detailed information about the commands used to investigate the problem.

Solution

[Table 52 on page 1943](#) provides commands for verifying the physical layer.

Table 52: Checklist for Verifying the Physical Layer

Tasks	Command or Action
“Verifying the Physical Layer” on page 1944	
1. Verify the LSP on page 1946	show mpls lsp extensive
2. Verify Router Connection on page 1948	ping host
3. Verify Interfaces on page 1949	show interfaces terse show configuration interfaces type-fpc/pic/port

Table 52: Checklist for Verifying the Physical Layer *(continued)*

Tasks	Command or Action
4. Take Appropriate Action on page 1950	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre>[edit interfaces <i>type-fpc/pic/port</i>] set family mpls show commit</pre>
5. Verify the LSP Again on page 1951	<pre>show mpls lsp extensive</pre>

Verifying the Physical Layer

Purpose

After you have configured the LSP, issued the **show mpls lsp extensive** command, and determined that there is an error, you can start investigating the problem at the physical layer of the network.

[Figure 150 on page 1945](#) illustrates the physical layer of the layered MPLS model.

Figure 150: Verifying the Physical Layer

BGP Layer	tracroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> tracroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
IGP and IP Layers Functioning	
OSPF Layer	IS-IS Layer
show ospf neighbor show configuration protocols ospf show ospf interface	show isis adjacency show configuration protocols isis show isis interface
IP Layer	IP Layer
show ospf neighbor extensive show interfaces terse	show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive "JUNOS Interfaces Operations Guide"
Physical Layer	show interfaces show interfaces terse ping <i>host</i>


g015543

With this layer, you must ensure that the routers are connected, and that the interfaces are up and configured correctly on the ingress, egress, and transit routers.

If the network is not functioning at this layer, the label-switched path (LSP) does not work as configured.

[Figure 151 on page 1946](#) illustrates the MPLS network and the problem described in this topic.

Figure 151: MPLS Network Broken at the Physical Layer



The network shown in [Figure 151 on page 1946](#) is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, traffic does not use the configured LSP. Instead traffic uses the alternate route from **R1** through **R2** to **R6**, and in the reverse direction, from **R6** through **R5** to **R1**.

When you become aware of a situation where an alternate route is used rather than the configured LSP, verify that the physical layer is functioning correctly. You might find that routers are not connected, or that interfaces are not up and configured correctly on the ingress, egress, or transit routers.

The cross shown in [Figure 151 on page 1946](#) indicates where the LSP is broken because of a configuration error on ingress router **R1**.

To check the physical layer, follow these steps:

1. [Verify the LSP | 1946](#)
2. [Verify Router Connection | 1948](#)
3. [Verify Interfaces | 1949](#)
4. [Take Appropriate Action | 1950](#)
5. [Verify the LSP Again | 1951](#)

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However, for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To determine whether the LSP is up, enter the following command from the ingress router:

```
user@ingress-router> show mpls lsp extensive
```

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.12.2 S 10.1.26.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.12.2 10.1.26.2
99 Sep 18 14:19:04 CSPF: computation result accepted
98 Sep 18 14:19:04 CSPF: link down/deleted
10.1.13.1(R1.00/10.0.0.1)->10.1.13.2(R3.00/10.0.0.3)
97 Sep 18 14:19:01 Record Route: 10.1.12.2 10.1.26.2
96 Sep 18 14:19:01 Up
95 Sep 18 14:19:01 Clear Call
94 Sep 18 14:19:01 CSPF: computation result accepted
93 Sep 18 14:19:01 MPLS label allocation failure
92 Sep 18 14:19:01 Down
91 Aug 17 12:22:52 Selected as active path
90 Aug 17 12:22:52 Record Route: 10.1.13.2 10.1.36.2
89 Aug 17 12:22:52 Up
[...Output truncated...]
Created: Sat Jul 10 18:18:44 2004
Total 1 displayed, Up 1, Down 0
```



```

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 144, Since: Tue Aug 17 12:23:14 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 67333 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.56.2 10.1.15.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output from ingress router **R1** shows that the LSP is using an alternate path rather than the configured path. The configured path for the LSP is **R1** through **R3** to **R6**, and for the reverse LSP, **R6** through **R3** to **R1**. The alternate path used by the LSP is **R1** through **R2** to **R6**, and for the reverse LSP, **R6** through **R5** to **R1**.

Verify Router Connection

Purpose

Confirm that the appropriate ingress, transit, and egress routers are functioning by examining whether the packets have been received and transmitted with 0% packet loss.

Action

To determine that the routers are connected, enter the following command from the ingress and transit routers:

```
user@host> ping host
```


Sample Output

```

user@R1> ping 10.0.0.3 count 3
PING 10.0.0.3 (10.0.0.3): 56 data bytes
64 bytes from 10.0.0.3: icmp_seq=0 ttl=254 time=0.859 ms
64 bytes from 10.0.0.3: icmp_seq=1 ttl=254 time=0.746 ms
64 bytes from 10.0.0.3: icmp_seq=2 ttl=254 time=0.776 ms

--- 10.0.0.3 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.746/0.794/0.859/0.048 ms

user@R3> ping 10.0.0.6 count 3
PING 10.0.0.6 (10.0.0.6): 56 data bytes
64 bytes from 10.0.0.6: icmp_seq=0 ttl=255 time=0.968 ms
64 bytes from 10.0.0.6: icmp_seq=1 ttl=255 time=3.221 ms
64 bytes from 10.0.0.6: icmp_seq=2 ttl=255 time=0.749 ms

--- 10.0.0.6 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.749/1.646/3.221/1.117 ms

```

Meaning

The sample output shows that ingress router **R1** is receiving packets from transit router **R3**, and that the transit router is receiving packets from the egress router. Therefore, the routers in the LSP are connected.

Verify Interfaces

Purpose

Confirm that the interfaces are configured correctly with the **family mpls** statement.

Action

To determine that the relevant interfaces are up and configured correctly, enter the following commands from the ingress, transit, and egress routers:

```

user@host> show interfaces terse
user@host> show configuration interfaces type-fpc/pic/port

```


Sample Output

```

user@R1> show interfaces so* terse
Interface          Admin Link Proto Local          Remote
so-0/0/0           up   up
so-0/0/0.0         up   up   inet  10.1.12.1/30
                               iso
                               mpls
so-0/0/1           up   up
so-0/0/1.0         up   up   inet  10.1.15.1/30
                               iso
                               mpls
so-0/0/2           up   up
so-0/0/2.0         up   up   inet  10.1.13.1/30
                               iso  <<< family mpls is missing
so-0/0/3           up   down

```

```

user@R1> show configuration interfaces so-0/0/2
unit 0 {
    family inet {
        address 10.1.13.1/30;
    }
    family iso; <<< family mpls is missing
}

```

Meaning

The sample output shows that interface **so-0/0/2.0** on the ingress router does not have the **family mpls** statement configured at the **[edit interfaces type-fpc/pic/port]** hierarchy level, indicating that the interface is incorrectly configured to support the LSP. The LSP is configured correctly at the **[edit protocols mpls]** hierarchy level.

The output from the transit and egress routers (not shown) shows that the interfaces on those routers are configured correctly.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the **family mpls** statement, which was missing, is included in the configuration of ingress router **R1**.

Solution

To correct the error in this example, enter the following commands:

```
[edit interfaces type-fpc/pic/port]  
user@R1# set family mpls  
user@R1# show  
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/2 unit 0]  
user@R1# set family mpls  
  
[edit interfaces so-0/0/2 unit 0]  
user@R1# show  
family inet {  
    address 10.1.13.1/30;  
}  
family iso;  
family mpls;  
  
[edit interfaces so-0/0/2 unit 0]  
user@R1# commit  
commit complete
```

Meaning

The sample output from ingress router **R1** shows that the **family mpls** statement is configured correctly for interface **so-0/0/2.0**, and that the LSP is now functioning as originally configured.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the physical layer has been resolved.

Action

To verify that the LSP is up and traversing the network as expected, enter the following command:

```
user@host> show mpls lsp extensive
```


Sample Output 1

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
    112 Sep 21 16:27:33 Record Route: 10.1.13.2 10.1.36.2
    111 Sep 21 16:27:33 Up
    110 Sep 21 16:27:33 CSPF: computation result accepted
    109 Sep 21 16:27:33 CSPF: link down/deleted
10.1.12.1(R1.00/10.0.0.1)->10.1.12.2(R2.00/10.0.0.2)
    108 Sep 21 16:27:33 CSPF: link down/deleted
10.1.15.1(R1.00/10.0.0.1)->10.1.15.2(R5.00/10.0.0.5)
  [Output truncated...]
  Created: Sat Jul 10 18:18:44 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 149, Since: Tue Sep 21 16:29:43 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 7 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

```



```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
[edit protocols mpls]
user@R1# show
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
```

Meaning

Sample Output 1 from ingress router **R1** shows that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Sample Output 2 from ingress router **R1** shows that the LSP is forced to take the intended path because MPLS is deactivated on **R1** interfaces **so-0/0/0.0** and **so-0/0/1.0**. If these interfaces were not deactivated, even though the configuration is now correct, the LSP would still traverse the network through the alternate path.

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However, for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To determine whether the LSP is up, enter the following command from the ingress router:


```
user@ingress-router> show mpls lsp extensive
```

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.12.2 S 10.1.26.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.12.2 10.1.26.2
    99 Sep 18 14:19:04 CSPF: computation result accepted
    98 Sep 18 14:19:04 CSPF: link down/deleted
    10.1.13.1(R1.00/10.0.0.1)->10.1.13.2(R3.00/10.0.0.3)
    97 Sep 18 14:19:01 Record Route: 10.1.12.2 10.1.26.2
    96 Sep 18 14:19:01 Up
    95 Sep 18 14:19:01 Clear Call
    94 Sep 18 14:19:01 CSPF: computation result accepted
    93 Sep 18 14:19:01 MPLS label allocation failure
    92 Sep 18 14:19:01 Down
    91 Aug 17 12:22:52 Selected as active path
    90 Aug 17 12:22:52 Record Route: 10.1.13.2 10.1.36.2
    89 Aug 17 12:22:52 Up
    [...Output truncated...]
    Created: Sat Jul 10 18:18:44 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
```



```

Resv style: 1 FF, Label in: 3, Label out: -
Time left: 144, Since: Tue Aug 17 12:23:14 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39024 protocol 0
PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 67333 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.56.2 10.1.15.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output from ingress router **R1** shows that the LSP is using an alternate path rather than the configured path. The configured path for the LSP is **R1** through **R3** to **R6**, and for the reverse LSP, **R6** through **R3** to **R1**. The alternate path used by the LSP is **R1** through **R2** to **R6**, and for the reverse LSP, **R6** through **R5** to **R1**.

Verify Router Connection

Purpose

Confirm that the appropriate ingress, transit, and egress routers are functioning by examining whether the packets have been received and transmitted with 0% packet loss.

Action

To determine that the routers are connected, enter the following command from the ingress and transit routers:

```
user@host> ping host
```

Sample Output

```

user@R1> ping 10.0.0.3 count 3
PING 10.0.0.3 (10.0.0.3): 56 data bytes
64 bytes from 10.0.0.3: icmp_seq=0 ttl=254 time=0.859 ms

```



```

64 bytes from 10.0.0.3: icmp_seq=1 ttl=254 time=0.746 ms
64 bytes from 10.0.0.3: icmp_seq=2 ttl=254 time=0.776 ms

--- 10.0.0.3 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.746/0.794/0.859/0.048 ms

user@R3> ping 10.0.0.6 count 3
PING 10.0.0.6 (10.0.0.6): 56 data bytes
64 bytes from 10.0.0.6: icmp_seq=0 ttl=255 time=0.968 ms
64 bytes from 10.0.0.6: icmp_seq=1 ttl=255 time=3.221 ms
64 bytes from 10.0.0.6: icmp_seq=2 ttl=255 time=0.749 ms

--- 10.0.0.6 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.749/1.646/3.221/1.117 ms

```

Meaning

The sample output shows that ingress router **R1** is receiving packets from transit router **R3**, and that the transit router is receiving packets from the egress router. Therefore, the routers in the LSP are connected.

Verify Interfaces

Purpose

Confirm that the interfaces are configured correctly with the **family mpls** statement.

Action

To determine that the relevant interfaces are up and configured correctly, enter the following commands from the ingress, transit, and egress routers:

```

user@host> show interfaces terse
user@host> show configuration interfaces type-fpc/pic/port

```

Sample Output

```

user@R1> show interfaces so* terse

```

Interface	Admin	Link	Proto	Local	Remote


```

so-0/0/0          up    up
so-0/0/0.0        up    up    inet  10.1.12.1/30
                               iso
                               mpls

so-0/0/1          up    up
so-0/0/1.0        up    up    inet  10.1.15.1/30
                               iso
                               mpls

so-0/0/2          up    up
so-0/0/2.0        up    up    inet  10.1.13.1/30
                               iso    <<< family mpls is missing

so-0/0/3          up    down

user@R1> show configuration interfaces so-0/0/2
unit 0 {
    family inet {
        address 10.1.13.1/30;
    }
    family iso; <<< family mpls is missing
}

```

Meaning

The sample output shows that interface **so-0/0/2.0** on the ingress router does not have the **family mpls** statement configured at the [edit interfaces type-fpc/pic/port] hierarchy level, indicating that the interface is incorrectly configured to support the LSP. The LSP is configured correctly at the [edit protocols mpls] hierarchy level.

The output from the transit and egress routers (not shown) shows that the interfaces on those routers are configured correctly.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the **family mpls** statement, which was missing, is included in the configuration of ingress router **R1**.

Solution

To correct the error in this example, enter the following commands:


```
[edit interfaces type-fpc/pic/port]
user@R1# set family mpls
user@R1# show
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/2 unit 0]
user@R1# set family mpls

[edit interfaces so-0/0/2 unit 0]
user@R1# show
family inet {
    address 10.1.13.1/30;
}
family iso;
family mpls;

[edit interfaces so-0/0/2 unit 0]
user@R1# commit
commit complete
```

Meaning

The sample output from ingress router **R1** shows that the **family mpls** statement is configured correctly for interface **so-0/0/2.0**, and that the LSP is now functioning as originally configured.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the physical layer has been resolved.

Action

To verify that the LSP is up and traversing the network as expected, enter the following command:

```
user@host> show mpls lsp extensive
```


Sample Output 1

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
    112 Sep 21 16:27:33 Record Route: 10.1.13.2 10.1.36.2
    111 Sep 21 16:27:33 Up
    110 Sep 21 16:27:33 CSPF: computation result accepted
    109 Sep 21 16:27:33 CSPF: link down/deleted
10.1.12.1(R1.00/10.0.0.1)->10.1.12.2(R2.00/10.0.0.2)
    108 Sep 21 16:27:33 CSPF: link down/deleted
10.1.15.1(R1.00/10.0.0.1)->10.1.15.2(R5.00/10.0.0.5)
  [Output truncated...]
  Created: Sat Jul 10 18:18:44 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 149, Since: Tue Sep 21 16:29:43 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 7 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

```



```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
[edit protocols mpls]
user@R1# show
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
```

Meaning

Sample Output 1 from ingress router **R1** shows that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Sample Output 2 from ingress router **R1** shows that the LSP is forced to take the intended path because MPLS is deactivated on **R1** interfaces **so-0/0/0.0** and **so-0/0/1.0**. If these interfaces were not deactivated, even though the configuration is now correct, the LSP would still traverse the network through the alternate path.

Checklist for Checking the Data Link Layer

Problem

Description: This checklist provides the steps and commands for investigating a problem at the data link layer of the Multiprotocol Label Switching (MPLS) network. The checklist provides links to an overview of the data link layer and more detailed information about the commands used to investigate the problem.

Solution

[Table 53 on page 1961](#) provides commands for checking the data link layer.

Table 53: Checklist for Checking the Data Link Layer

Tasks	Command or Action
“Checking the Data Link Layer” on page 1961	
1. Verify the LSP on page 1963	show mpls lsp extensive
2. Verify Interfaces on page 1964	show interfaces <i>type-fpc/pic/port</i> extensive show interfaces <i>type-fpc/pic/port</i>
3. Take Appropriate Action on page 1969	The following sequence of commands addresses the specific problem described in this topic: [edit interfaces <i>type-fpc/pic/port</i>] show delete encapsulation show commit
4. Verify the LSP Again on page 1970	show mpls lsp extensive

Checking the Data Link Layer

Purpose

After you have configured the label-switched path (LSP), issued the **show mpls lsp extensive** command, and determined that there is an error, you might find that the error is not in the physical layer. Continue investigating the problem at the data link layer of the network.

[Figure 152 on page 1962](#) illustrates the data link layer of the layered MPLS model.


Figure 152: Checking the Data Link Layer

BGP Layer	tracroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> tracroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
↩ IGP and IP Layers Functioning ↪	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive "JUNOS Interfaces Operations Guide"
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

With this layer, you must check the encapsulation mode, for example, Point-to-Point Protocol (PPP) or Cisco High-Level Data Link Control (HDLC); PPP options, for example, header encapsulation; frame check sequence (FCS) size; and whether keepalive frames are enabled or disabled. Also, check the ingress, egress, and transit routers.

Figure 153 on page 1962 illustrates the MPLS network used in this topic.

Figure 153: MPLS Network Broken at the Data Link Layer



The network shown in [Figure 153 on page 1962](#) is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, the LSP is down without a path in either direction, from **R1** to **R6** or from **R6** to **R1**.

When you verify that the data link layer is not functioning correctly, you might find a mismatch with PPP or Cisco HDLC encapsulation, PPP options, or keepalive frames.

The cross shown in [Figure 153 on page 1962](#) indicates where the LSP is broken because of a configuration error on ingress router **R1** that prevents the LSP from traversing the network as expected.

To check the data link layer, follow these steps:

1. [Verify the LSP | 1963](#)
2. [Verify Interfaces | 1964](#)
3. [Take Appropriate Action | 1969](#)
4. [Verify the LSP Again | 1970](#)

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To determine whether the LSP is up, enter the following command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
```



```

From: 10.0.0.1 , State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
ActivePath: (none)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary                               State: Dn
    Will be enqueued for recomputation in 15 second(s).
140 Sep 30 12:01:12 CSPF failed: no route toward 10.0.0.6[26 times]
139 Sep 30 11:48:57 Deselected as active
138 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
137 Sep 30 11:48:56 Clear Call
136 Sep 30 11:48:56 CSPF: link down/deleted
10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)
135 Sep 30 11:48:56 ResvTear received
134 Sep 30 11:48:56 Down
133 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
132 Sep 30 11:48:56 10.1.13.2: No Route toward dest
[...Output truncated...]
Created: Sat Jul 10 18:18:44 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output from ingress router **R1** shows the LSPs within which it participates. The ingress LSP is down, without a path from **R1** to **R6**. Because a reverse LSP is configured in the network shown in [Figure 153 on page 1962](#), we would expect an egress LSP session to be up. However, **R1** does not have any egress LSPs, indicating that the LSP from **R6** to **R1** is not functioning.

Verify Interfaces

Purpose

From your network topology, determine the adjacent interfaces through which the LSP is meant to traverse, and examine the output for the encapsulation type, PPP options, FCS size, and whether keepalive frames are enabled or disabled

NOTE: Before you proceed with this step, check the physical layer to ensure that the problem is not in the physical layer.

Action

To verify the functioning of adjacent interfaces, enter the following commands from the relevant routers:

```
user@host> show interfaces type-fpc/pic/port extensive
user@host> show interfaces type-fpc/pic/port
```

Sample Output 1

```
user@R6> show interfaces so-0/0/3 extensive
Physical interface: so-0/0/3, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 27, Generation: 14
    Link-level type: Cisco-HDLC , MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
    Loopback: None,
    FCS: 16 , Payload scrambler: Enabled
    Device flags      : Present Running
    Interface flags: Link-Layer-Down   Point-To-Point SNMP-Traps 16384
    Link flags       : Keepalives
    Hold-times       : Up 0 ms, Down 0 ms
    Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Keepalive statistics:
      Input  : 0 (last seen: never)
      Output: 357 (last sent 00:00:04 ago)
    CoS queues       : 4 supported
    Last flapped     : 2004-07-21 16:03:49 PDT (10w0d 07:01 ago)
    Statistics last cleared: Never
    Traffic statistics:
      Input  bytes   :          203368873          0 bps
      Output bytes   :          186714992         88 bps
      Input  packets :          3641808          0 pps
      Output packets :          3297569          0 pps
    Input errors:
      Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Bucket drops: 0,
      Policed discards: 1770, L3 incompletes: 0, L2 channel errors: 0, L2 mismatch
```



```

timeouts: 0,
    HS link CRC errors: 0, HS link FIFO overflows: 0
Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0, HS link FIFO
underflows: 0,
    MTU errors: 0
Queue counters:      Queued packets  Transmitted packets      Dropped packets
0 best-effort        197012                197012                    0
1 expedited-fo        0                      0                        0
2 assured-forw        0                      0                        0
3 network-cont       3100557                3100557                    0

SONET alarms : None
SONET defects : None
SONET PHY:           Seconds          Count    State
    PLL Lock          0              0    OK
    PHY Light          0              0    OK
SONET section:
    BIP-B1             0              0
    SEF                 1              3    OK
    LOS                 1              1    OK
    LOF                 1              1    OK
    ES-S                1
    SES-S               1
    SEFS-S              1
SONET line:
    BIP-B2             0              0
    REI-L              0              0
    RDI-L              0              0    OK
    AIS-L              0              0    OK
    BERR-SF            0              0    OK
    BERR-SD            0              0    OK
    ES-L               1
    SES-L              1
    UAS-L              0
    ES-LFE             0
    SES-LFE            0
    UAS-LFE            0
SONET path:
    BIP-B3             0              0
    REI-P              0              0
    LOP-P              0              0    OK
    AIS-P              0              0    OK

```


RDI-P	0	0	OK
UNEQ-P	0	0	OK
PLM-P	0	0	OK
ES-P	1		
SES-P	1		
UAS-P	0		
ES-PFE	0		
SES-PFE	0		
UAS-PFE	0		

Received SONET overhead:

F1	: 0x00, J0	: 0x00, K1	: 0x00, K2	: 0x00
S1	: 0x00, C2	: 0xcf, C2(cmp)	: 0xcf, F2	: 0x00
Z3	: 0x00, Z4	: 0x00, S1(cmp)	: 0x00	

Transmitted SONET overhead:

F1	: 0x00, J0	: 0x01, K1	: 0x00, K2	: 0x00
S1	: 0x00, C2	: 0xcf, F2	: 0x00, Z3	: 0x00
Z4	: 0x00			

Received path trace: R3 so-0/0/3

52 33 20 73 6f 2d 30 2f 30 2f 33 00 00 00 00 00	R3 so-0/0/3.. ...
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a

Transmitted path trace: R6 so-0/0/3

52 36 20 73 6f 2d 30 2f 30 2f 33 00 00 00 00 00	R6 so-0/0/3
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

HDLC configuration:

Policing bucket: Disabled
Shaping bucket : Disabled
Giant threshold: 4484, Runt threshold: 3

Packet Forwarding Engine configuration:

Destination slot: 0, PLP byte: 1 (0x00)						
CoS transmit queue	Bandwidth		Buffer		Priority	Limit
	%	bps	%	bytes		
0 best-effort	95	147744000	95	0	low	none
3 network-control	5	7776000	5	0	low	none

Logical interface so-0/0/3.0 (Index 71) (SNMP ifIndex 28) (Generation 16)

Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: Cisco-HDLC

Traffic statistics:

Input bytes :	406737746
Output bytes :	186714992


```

    Input  packets:          7283616
    Output packets:          3297569
Local statistics:
    Input  bytes   :          203368873
    Output bytes   :          186714992
    Input  packets:          3641808
    Output packets:          3297569
Transit statistics:
    Input  bytes   :          203368873          0 bps
    Output bytes   :                   0          0 bps
    Input  packets:          3641808          0 pps
    Output packets:                   0          0 pps
Protocol inet, MTU: 4470, Generation: 46, Route table: 0
  Flags: None
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.36.0/30, Local: 10.1.36.2, Broadcast: 10.1.36.3, Generation: 38
Protocol iso, MTU: 4469, Generation: 47, Route table: 0
  Flags: None
Protocol mpls, MTU: 4458, Generation: 48, Route table: 0
  Flags: None

```

Sample Output 2

```

user@R3> show interfaces so-0/0/3
Physical interface: so-0/0/3, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 24
    Link-level type: PPP , MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16 ,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
    Link flags   : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 736827 (00:00:03 ago), Output: 736972 (00:00:05 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  CoS queues      : 4 supported
  Last flapped    : 2004-07-21 16:08:01 PDT (10w5d 19:57 ago)
  Input rate      : 40 bps (0 pps)

```



```

Output rate      : 48 bps (0 pps)
SONET alarms     : None
SONET defects    : None

```

Logical interface so-0/0/3.0 (Index 70) (SNMP ifIndex 51)

Flags: Point-To-Point SNMP-Traps **Encapsulation: PPP**

Protocol inet, MTU: 4470

Flags: None

Addresses, Flags: Is-Preferred Is-Primary

Destination: 10.1.36.0/30, Local: 10.1.36.1, Broadcast: 10.1.36.3

Protocol iso, MTU: 4470

Flags: None

Protocol mpls, MTU: 4458

Flags: None

Meaning

Sample Output 1 from egress router **R6** shows that there are no SONET alarms or defects (**none**), the states are all **OK**, and the path trace shows the distant end (**R3 so-0.0.0**), indicating that the physical link is up. However, the logical link is down, and the link-level type is Cisco HDLC.

Sample Output 2 from transit router **R3** shows that the link-level type is PPP, indicating that the encapsulation types are mismatched, resulting in the LSP going down.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the encapsulation types are mismatched.

Solution

To correct the error in this example, enter the following commands:

```

[edit interfaces so-0/0/3]
user@R1# show
user@R1# delete encapsulation
user@R1# show
user@R1# commit

```


Sample Output

```
[edit interfaces so-0/0/3]
user@R6# show
encapsulation cisco-hdlc;
unit 0 {
    family inet {
        address 10.1.36.2/30;
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/3]
user@R6# delete encapsulation

[edit interfaces so-0/0/3]
user@R6# show
unit 0 {
    family inet {
        address 10.1.36.2/30;
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/3]
user@R6# commit
commit complete
```

Meaning

The sample output from egress router **R6** shows that the Cisco HDLC was incorrectly configured on interface **so-0/0/3** which prevented the LSP from using the intended path. The problem was corrected when the **encapsulation** statement was deleted and the configuration committed.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the data link layer has been resolved.

Action

From the ingress, egress, and transit routers, verify that the LSP is up and traversing the network as expected:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1 , State: Up,  ActiveRoute: 1 , LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
  10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
145 Sep 30 12:25:01 Selected as active path
144 Sep 30 12:25:01 Record Route:  10.1.13.2 10.1.36.2
143 Sep 30 12:25:01 Up
142 Sep 30 12:25:01 Originate Call
141 Sep 30 12:25:01 CSPF: computation result accepted
140 Sep 30 12:24:32 CSPF failed: no route toward 10.0.0.6[74 times]
139 Sep 30 11:48:57 Deselected as active
138 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
137 Sep 30 11:48:56 Clear Call
136 Sep 30 11:48:56 CSPF: link down/deleted
10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)
  [...Output truncated...]
  Created: Sat Jul 10 18:18:43 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6 , LSPstate: Up, ActiveRoute: 0
```



```

LSPname: R6-to-R1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 134, Since: Thu Sep 30 12:24:56 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 6 receiver 39024 protocol 0
PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 7 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient

Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 2

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.1.36.1 10.1.13.1
50 Sep 30 12:24:12 Selected as active path
49 Sep 30 12:24:12 Record Route: 10.1.36.1 10.1.13.1
48 Sep 30 12:24:12 Up
47 Sep 30 12:24:12 Originate Call
46 Sep 30 12:24:12 CSPF: computation result accepted
45 Sep 30 12:23:43 CSPF failed: no route toward 10.0.0.1[73 times]
44 Sep 30 11:48:12 Deselected as active

```



```

    43 Sep 30 11:48:12 CSPF failed: no route toward 10.0.0.1
    42 Sep 30 11:48:12 CSPF: link down/deleted
10.1.36.2(R6.00/10.0.0.6)->10.1.36.1(R3.00/10.0.0.3)
    [...Output truncated...]
    Created: Tue Aug 17 12:18:34 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

```

10.0.0.6

```

From: 10.0.0.1 , LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 159, Since: Thu Sep 30 12:24:16 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 19 receiver 44251 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient

Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 3

```

user@R3> show mpls lsp extensive

Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1

```



```

From: 10.0.0.6 , LSPstate: Up, ActiveRoute: 1
LSPname: R6-to-R1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100176, Label out: 3
Time left: 143, Since: Thu Sep 30 12:21:25 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 6 receiver 39024 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 9 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 9 pkts
Explct route: 10.1.13.1

Record route: 10.1.36.2 <self> 10.1.13.1

```

10.0.0.6

```

From: 10.0.0.1 , LSPstate: Up, ActiveRoute: 1
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100192, Label out: 3
Time left: 148, Since: Thu Sep 30 12:21:30 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 19 receiver 44251 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 9 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 9 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 9 pkts
Explct route: 10.1.36.2

Record route: 10.1.13.1 <self> 10.1.36.2

```

Total 2 displayed, Up 2, Down 0

Sample Output 4

```

user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;

```



```

interface so-0/0/2.0;

user@R6> show configuration protocols mpls
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
interface so-0/0/3.0;

user@R3> show configuration protocols mpls
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;

```

Meaning

Sample Outputs 1 and 2 from ingress router **R1** and egress router **R6**, respectively, show that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Sample Output 3 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**.

Sample Output 4 shows the interfaces that were deactivated on the ingress, egress, and transit routers, forcing the LSP to take the intended path. If these interfaces were not deactivated, even though the configuration is now correct, the LSP would still traverse the network through the alternate path.

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To determine whether the LSP is up, enter the following command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1 , State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 15 second(s).
    140 Sep 30 12:01:12 CSPF failed: no route toward 10.0.0.6[26 times]
    139 Sep 30 11:48:57 Deselected as active
    138 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
    137 Sep 30 11:48:56 Clear Call
    136 Sep 30 11:48:56 CSPF: link down/deleted
    10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)
    135 Sep 30 11:48:56 ResvTear received
    134 Sep 30 11:48:56 Down
    133 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6
    132 Sep 30 11:48:56 10.1.13.2: No Route toward dest
    [...Output truncated...]
    Created: Sat Jul 10 18:18:44 2004
  Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

The sample output from ingress router **R1** shows the LSPs within which it participates. The ingress LSP is down, without a path from **R1** to **R6**. Because a reverse LSP is configured in the network shown in [Figure](#)

153 on page 1962, we would expect an egress LSP session to be up. However, **R1** does not have any egress LSPs, indicating that the LSP from **R6** to **R1** is not functioning.

Verify Interfaces

Purpose

From your network topology, determine the adjacent interfaces through which the LSP is meant to traverse, and examine the output for the encapsulation type, PPP options, FCS size, and whether keepalive frames are enabled or disabled

NOTE: Before you proceed with this step, check the physical layer to ensure that the problem is not in the physical layer.

Action

To verify the functioning of adjacent interfaces, enter the following commands from the relevant routers:

```
user@host> show interfaces type-fpc/pic/port extensive
user@host> show interfaces type-fpc/pic/port
```

Sample Output 1

```
user@R6> show interfaces so-0/0/3 extensive
Physical interface: so-0/0/3, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 27, Generation: 14
    Link-level type: Cisco-HDLC , MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
    Loopback: None,
    FCS: 16 , Payload scrambler: Enabled
    Device flags   : Present Running
    Interface flags: Link-Layer-Down  Point-To-Point SNMP-Traps 16384
    Link flags    : Keepalives
    Hold-times    : Up 0 ms, Down 0 ms
    Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
    Keepalive statistics:
      Input : 0 (last seen: never)
      Output: 357 (last sent 00:00:04 ago)
```


CoS queues : 4 supported

Last flapped : 2004-07-21 16:03:49 PDT (10w0d 07:01 ago)

Statistics last cleared: Never

Traffic statistics:

Input bytes :	203368873	0 bps
Output bytes :	186714992	88 bps
Input packets:	3641808	0 pps
Output packets:	3297569	0 pps

Input errors:

Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Bucket drops: 0,

Policed discards: 1770, L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,

HS link CRC errors: 0, HS link FIFO overflows: 0

Output errors:

Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0, HS link FIFO underflows: 0,

MTU errors: 0

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	197012	197012	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	3100557	3100557	0

SONET alarms : None

SONET defects : None

SONET PHY:	Seconds	Count	State
PLL Lock	0	0	OK
PHY Light	0	0	OK
SONET section:			
BIP-B1	0	0	
SEF	1	3	OK
LOS	1	1	OK
LOF	1	1	OK
ES-S	1		
SES-S	1		
SEFS-S	1		
SONET line:			
BIP-B2	0	0	
REI-L	0	0	
RDI-L	0	0	OK
AIS-L	0	0	OK
BERR-SF	0	0	OK
BERR-SD	0	0	OK


```

ES-L                1
SES-L               1
UAS-L               0
ES-LFE              0
SES-LFE              0
UAS-LFE              0
SONET path:
  BIP-B3             0          0
  REI-P              0          0
  LOP-P              0          0    OK
  AIS-P              0          0    OK
  RDI-P              0          0    OK
  UNEQ-P             0          0    OK
  PLM-P              0          0    OK
  ES-P               1
  SES-P               1
  UAS-P               0
  ES-PFE              0
  SES-PFE              0
  UAS-PFE              0
Received SONET overhead:
  F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, C2(cmp) : 0xcf, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00
Received path trace: R3 so-0/0/3
52 33 20 73 6f 2d 30 2f 30 2f 33 00 00 00 00 00    R3 so-0/0/3.. ...
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00    .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00    .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a    .....
Transmitted path trace: R6 so-0/0/3
52 36 20 73 6f 2d 30 2f 30 2f 33 00 00 00 00 00    R6 so-0/0/3 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00    .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00    .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00    .....
HDLC configuration:
  Policing bucket: Disabled
  Shaping bucket : Disabled
  Giant threshold: 4484, Runt threshold: 3
Packet Forwarding Engine configuration:
```



```

Destination slot: 0, PLP byte: 1 (0x00)
CoS transmit queue      Bandwidth      Buffer Priority  Limit
                        %      bps      %      bytes
0 best-effort           95      147744000  95          0      low  none
3 network-control       5       7776000   5          0      low  none

```

Logical interface so-0/0/3.0 (Index 71) (SNMP ifIndex 28) (Generation 16)

Flags: Device-Down Point-To-Point SNMP-Traps **Encapsulation: Cisco-HDLC**

Traffic statistics:

```

Input  bytes :      406737746
Output bytes :      186714992
Input  packets:      7283616
Output packets:      3297569

```

Local statistics:

```

Input  bytes :      203368873
Output bytes :      186714992
Input  packets:      3641808
Output packets:      3297569

```

Transit statistics:

```

Input  bytes :      203368873      0 bps
Output bytes :           0      0 bps
Input  packets:      3641808      0 pps
Output packets:           0      0 pps

```

Protocol inet, MTU: 4470, Generation: 46, Route table: 0

Flags: None

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

Destination: 10.1.36.0/30, Local: 10.1.36.2, Broadcast: 10.1.36.3, Generation: 38

Protocol iso, MTU: 4469, Generation: 47, Route table: 0

Flags: None

Protocol mpls, MTU: 4458, Generation: 48, Route table: 0

Flags: None

Sample Output 2

```

user@R3> show interfaces so-0/0/3
Physical interface: so-0/0/3, Enabled, Physical link is Up
Interface index: 131, SNMP ifIndex: 24
Link-level type: PPP , MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16 ,
Payload scrambler: Enabled

```



```

Device flags      : Present Running
Interface flags: Point-To-Point SNMP-Traps
  Link flags      : Keepalives
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive: Input: 736827 (00:00:03 ago), Output: 736972 (00:00:05 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Opened, mpls: Opened
CHAP state: Not-configured
CoS queues       : 4 supported
Last flapped     : 2004-07-21 16:08:01 PDT (10w5d 19:57 ago)
Input rate       : 40 bps (0 pps)
Output rate      : 48 bps (0 pps)
SONET alarms     : None
SONET defects    : None

```

Logical interface so-0/0/3.0 (Index 70) (SNMP ifIndex 51)

```

Flags: Point-To-Point SNMP-Traps  Encapsulation: PPP
Protocol inet, MTU: 4470
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.1.36.0/30, Local: 10.1.36.1, Broadcast: 10.1.36.3
Protocol iso, MTU: 4470
  Flags: None
Protocol mpls, MTU: 4458
  Flags: None

```

Meaning

Sample Output 1 from egress router **R6** shows that there are no SONET alarms or defects (**none**), the states are all **OK**, and the path trace shows the distant end (**R3 so-0.0.0**), indicating that the physical link is up. However, the logical link is down, and the link-level type is Cisco HDLC.

Sample Output 2 from transit router **R3** shows that the link-level type is PPP, indicating that the encapsulation types are mismatched, resulting in the LSP going down.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the encapsulation types are mismatched.

Solution

To correct the error in this example, enter the following commands:

```
[edit interfaces so-0/0/3]
user@R1# show
user@R1# delete encapsulation
user@R1# show
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/3]
user@R6# show
encapsulation cisco-hdlc;
unit 0 {
    family inet {
        address 10.1.36.2/30;
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/3]
user@R6# delete encapsulation

[edit interfaces so-0/0/3]
user@R6# show
unit 0 {
    family inet {
        address 10.1.36.2/30;
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/3]
user@R6# commit
commit complete
```

Meaning

The sample output from egress router **R6** shows that the Cisco HDLC was incorrectly configured on interface **so-0/0/3** which prevented the LSP from using the intended path. The problem was corrected when the **encapsulation** statement was deleted and the configuration committed.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the data link layer has been resolved.

Action

From the ingress, egress, and transit routers, verify that the LSP is up and traversing the network as expected:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
```

```
Ingress LSP: 1 sessions
```

10.0.0.6

From: 10.0.0.1 , State: Up, **ActiveRoute: 1** , **LSPname: R1-to-R6**

ActivePath: (primary)

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

*Primary State: Up

Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)

10.1.13.2 S 10.1.36.2 S

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.1.13.2 10.1.36.2

145 Sep 30 12:25:01 Selected as active path

144 Sep 30 12:25:01 Record Route: 10.1.13.2 10.1.36.2

143 Sep 30 12:25:01 Up

142 Sep 30 12:25:01 Originate Call

141 Sep 30 12:25:01 CSPF: computation result accepted

140 Sep 30 12:24:32 CSPF failed: no route toward 10.0.0.6[74 times]

139 Sep 30 11:48:57 Deselected as active

138 Sep 30 11:48:56 CSPF failed: no route toward 10.0.0.6

137 Sep 30 11:48:56 Clear Call

136 Sep 30 11:48:56 CSPF: link down/deleted

10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)

[...Output truncated...]


```

Created: Sat Jul 10 18:18:43 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6 , LSPstate: Up, ActiveRoute: 0
LSPname: R6-to-R1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 134, Since: Thu Sep 30 12:24:56 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 6 receiver 39024 protocol 0
PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 7 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient

Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 2

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

```



```

10.1.36.1 10.1.13.1
50 Sep 30 12:24:12 Selected as active path
49 Sep 30 12:24:12 Record Route: 10.1.36.1 10.1.13.1
48 Sep 30 12:24:12 Up
47 Sep 30 12:24:12 Originate Call
46 Sep 30 12:24:12 CSPF: computation result accepted
45 Sep 30 12:23:43 CSPF failed: no route toward 10.0.0.1[73 times]
44 Sep 30 11:48:12 Deselected as active
43 Sep 30 11:48:12 CSPF failed: no route toward 10.0.0.1
42 Sep 30 11:48:12 CSPF: link down/deleted
10.1.36.2(R6.00/10.0.0.6)->10.1.36.1(R3.00/10.0.0.3)
[...Output truncated...]
Created: Tue Aug 17 12:18:34 2004
Total 1 displayed, Up 1, Down 0

```

Egress LSP: 1 sessions

10.0.0.6

```

From: 10.0.0.1 , LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 159, Since: Thu Sep 30 12:24:16 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 19 receiver 44251 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

```

```

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 3

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions

```


Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1

From: 10.0.0.6 , LSPstate: Up, ActiveRoute: 1
 LSPname: R6-to-R1, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100176, Label out: 3
 Time left: 143, Since: Thu Sep 30 12:21:25 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 6 receiver 39024 protocol 0
 PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.13.1 (so-0/0/2.0) 9 pkts
 RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 9 pkts
 Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6

From: 10.0.0.1 , LSPstate: Up, ActiveRoute: 1
 LSPname: R1-to-R6, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100192, Label out: 3
 Time left: 148, Since: Thu Sep 30 12:21:30 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 19 receiver 44251 protocol 0
 PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 9 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.36.2 (so-0/0/3.0) 9 pkts
 RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 9 pkts
 Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2

Total 2 displayed, Up 2, Down 0

Sample Output 4

```

user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;

user@R6> show configuration protocols mpls
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
interface so-0/0/3.0;

user@R3> show configuration protocols mpls
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;

```

Meaning

Sample Outputs 1 and 2 from ingress router **R1** and egress router **R6**, respectively, show that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Sample Output 3 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**.

Sample Output 4 shows the interfaces that were deactivated on the ingress, egress, and transit routers, forcing the LSP to take the intended path. If these interfaces were not deactivated, even though the configuration is now correct, the LSP would still traverse the network through the alternate path.

Checklist for Verifying the IP and IGP Layers

Problem

Description: This checklist provides the steps and commands for investigating a problem at the Internet Protocol (IP) and interior gateway protocol (IGP) layers of the layered Multiprotocol Label Switching (MPLS) model. The checklist provides links to an overview of the IP and IGP layers and more detailed information about the commands used to investigate the problem.

Solution

Table 54 on page 1988 provides commands for verifying the IP and IGP layers.

Table 54: Checklist for Verifying the IP and IGP Layers

Tasks	Command or Action
“Verifying the IP and IGP Layers” on page 1990	
“Verifying the IP Layer” on page 1992	
1. Verify the LSP on page 1993	<code>show mpls lsp extensive</code>
2. Verify IP Addressing on page 1994	<code>show interfaces terse</code>
3. Verify Neighbors or Adjacencies at the IP Layer on page 1996	<code>show ospf neighbor extensive</code> <code>show isis adjacency extensive</code>
4. Take Appropriate Action on page 2000	The following sequence of commands addresses the specific problem described in this topic: [edit interfaces so-0/0/2] <code>show</code> <code>rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30</code> <code>show</code> <code>commit</code>
5. Verify the LSP Again on page 2002	<code>show mpls lsp extensive</code>
“Verifying the OSPF Protocol” on page 2019	
1. Verify the LSP on page 2020	<code>show mpls lsp extensive</code>
2. Verify OSPF Interfaces on page 2025	<code>show ospf interface</code>
3. Verify OSPF Neighbors on page 2027	<code>show ospf neighbor</code>

Table 54: Checklist for Verifying the IP and IGP Layers *(continued)*

Tasks	Command or Action
4. <i>Verifying an OSPF Configuration</i>	show configuration protocols ospf
5. <i>Taking Appropriate Action for Resolving the Network Problem</i>	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre> [edit] edit protocols ospf area 0.0.0.0 [edit protocols ospf area 0.0.0.0] set interface lo0 set interface lo0 passive up [edit protocols ospf] set traffic-engineering show commit </pre>
6. Verify the LSP Again on page 2031	show mpls lsp extensive
Verifying the IS-IS Protocol	
1. Verify the LSP on page 2046	show mpls lsp extensive
2. Verify IS-IS Adjacencies and Interfaces on page 2048	show isis adjacency show isis interface
3. Verify the IS-IS Configuration on page 2050	show configuration protocols isis
4. <i>Taking Appropriate Action for Resolving the Network Problem</i>	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre> edit [edit] edit protocols isis [edit protocols isis] show delete level 2 set level 1 disable show commit run show isis adjacency </pre>

Table 54: Checklist for Verifying the IP and IGP Layers *(continued)*

Tasks	Command or Action
5. Verify the LSP Again on page 2051	<code>show mpls lsp extensive</code>

Verifying the IP and IGP Layers

Problem

Description: After you have configured the label-switched path (LSP), issued the `show mpls lsp extensive` command, and determined that there is an error, you might find that the error is not in the physical or data link layers. Continue investigating the problem at the IP and IGP layers of the network.

[Figure 154 on page 1990](#) illustrates the IP and IGP layers of the layered MPLS model.

Figure 154: IP and IGP Layers

BGP Layer	<traceroute <i="">host-name show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i> </traceroute>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
<div> <div>IGP and IP Layers Functioning</div> <div> <div> <div> <div> <div>OSPF Layer</div> <div>show ospf neighbor</div> <div>show configuration protocols ospf</div> <div>show ospf interface</div> </div> <div> <div>IS-IS Layer</div> <div>show isis adjacency</div> <div>show configuration protocols isis</div> <div>show isis interface</div> </div> </div> <div> <div> <div>IP Layer</div> <div>show ospf neighbor extensive</div> <div>show interfaces terse</div> </div> <div> <div>IP Layer</div> <div>show isis adjacency extensive</div> <div>show interfaces terse</div> </div> </div> </div> </div></div>	
Data Link Layer	show interfaces extensive <i>"JUNOS Interfaces Operations Guide"</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

g015545

Solution

At the IP and IGP layers, you must check the following:

- Interfaces have correct IP addressing, and the IGP neighbors or adjacencies are established.

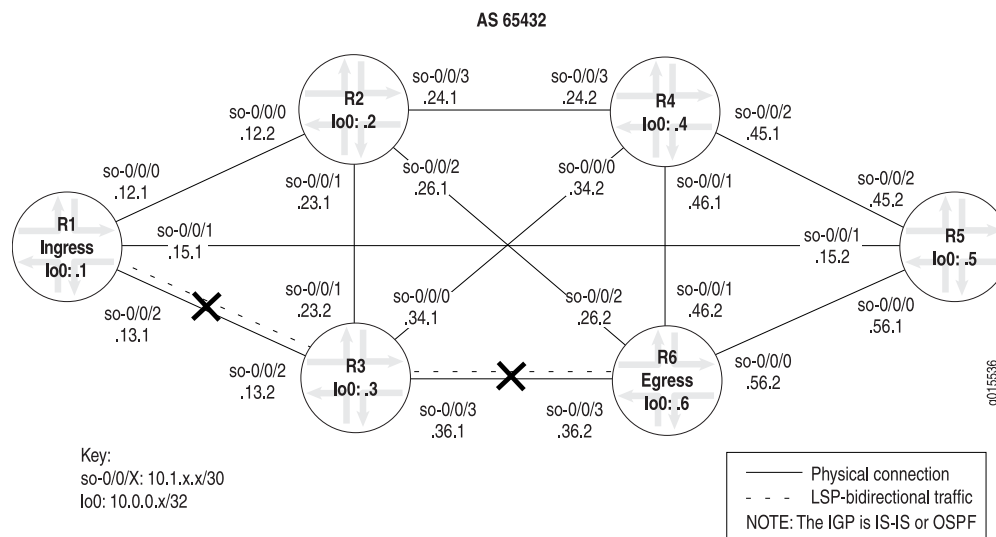
- Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) protocols are configured and running correctly.
- If the OSPF protocol is configured, check the IP layer first, then the OSPF configuration, making sure that the protocol, interfaces, and traffic engineering are configured correctly.
- If the IS-IS protocol is configured, it doesn't matter whether you check IS-IS or IP first because both protocols are independent of each other. Verify that IS-IS adjacencies are up, and that the interfaces and IS-IS protocol are configured correctly.

NOTE: The IS-IS protocol has traffic engineering enabled by default.

If the network is not functioning at the IP or IGP layers, the LSP does not work as configured.

Figure 155 on page 1991 illustrates the MPLS network used in this topic.

Figure 155: MPLS Network Broken at the IP and IGP Layers



The network shown in Figure 155 on page 1991 is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6**, through **R3**, to **R1**, creating bidirectional traffic. The crosses in Figure 155 on page 1991 indicate where the LSP is not working because of the following problems at the IP and IGP layer:

- An IP address is configured incorrectly on the ingress router (**R1**).
- The OSPF protocol is configured with a router ID (RID) but without the loopback (**lo0**) interface, and traffic engineering is missing from the transit router (**R3**).

- Levels in the IS-IS network are mismatched.

RELATED DOCUMENTATION

To check the IP and IGP layers, follow these steps:

[Verifying the IP Layer | 1992](#)

[Verifying the OSPF Protocol | 2019](#)

Verifying the IS-IS Protocol

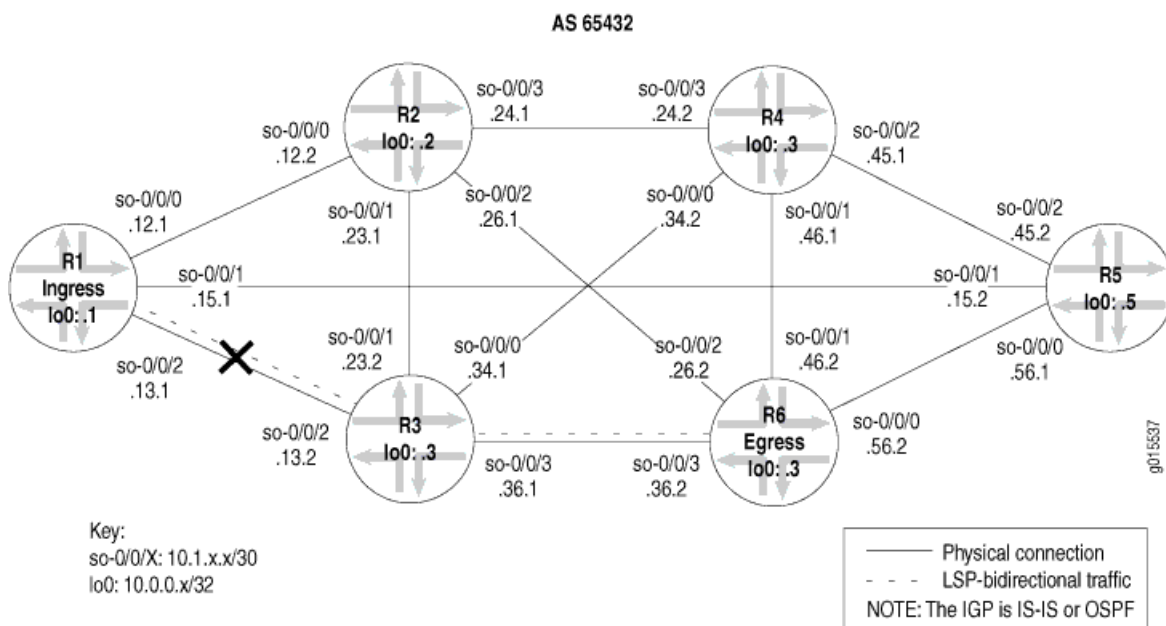
Verifying the IP Layer

Purpose

You can check the IP layer before or after you check the interior gateway protocol (IGP) layer, depending on whether you have OSPF or IS-IS configured as the IGP. If your MPLS network is configured with OSPF as the IGP, you must first verify the IP layer, checking that the interfaces have correct IP addressing and that the OSPF neighbors are established before you check the OSPF layer.

If you have IS-IS configured as the IGP in your MPLS network, you can verify either the IP layer or the IS-IS protocol layer first. The order in which you check the IP or IS-IS layer does not affect the results.

Figure 156: MPLS Network Broken at the IP Layer



The cross in [Figure 156 on page 1992](#) indicates where the LSP is broken because of the incorrect configuration of an IP address on ingress router **R1**.

1. [Verify the LSP | 1993](#)
2. [Verify IP Addressing | 1994](#)
3. [Verify Neighbors or Adjacencies at the IP Layer | 1996](#)
4. [Take Appropriate Action | 2000](#)
5. [Verify the LSP Again | 2002](#)

Verify the LSP

Purpose

After configuring the LSP, you must verify that the LSP is up. LSPs can be ingress, transit, or egress. Use the **show mpls lsp** command for quick verification of the LSP state, with the **extensive** option (**show mpls lsp extensive**) as a follow-up if the LSP is down. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To verify that the LSP is up, enter the following command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 25 second(s).
  44 Oct 15 16:56:11 CSPF failed:  no route toward 10.0.0.6 [2685 times]
  43 Oct 14 19:07:09 Clear Call
  42 Oct 14 19:06:56 Deselected as active
  41 Oct 14 19:06:56 10.1.12.1: MPLS label allocation failure
```



```

40 Oct 14 19:06:56 Down
39 Oct 14 18:43:43 Selected as active path
38 Oct 14 18:43:43 Record Route:  10.1.13.2 10.1.36.2
37 Oct 14 18:43:43 Up
[...Output truncated...]
Created: Thu Oct 14 16:04:33 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed , Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed , Up 0, Down 0

```

Meaning

The sample output from ingress router **R1** shows that an MPLS label allocation failure occurred and the Constrained Shortest Path First (CSPF) algorithm failed, resulting in no route to destination **10.0.0.6** on **R6**.

Verify IP Addressing

Purpose

When you investigate the IP layer, you verify that interfaces have correct IP addressing, and that OSPF neighbors or IS-IS adjacencies are established. In this example, an IP address is configured incorrectly on the ingress router (**R1**).

Action

To verify IP addressing, enter the following command from the ingress, transit, and egress routers:

```
user@host> show interfaces terse
```

Sample Output

```

user@R1> show interfaces terse
Interface          Admin Link Proto Local                Remote
so-0/0/0           up    up
so-0/0/0.0         up    up   inet 10.1.12.1/30
                   iso

```



```

mpls
so-0/0/1          up    up
so-0/0/1.0        up    up    inet  10.1.15.1/30
                               iso
                               mpls
so-0/0/2          up    up
so-0/0/2.0        up    up    inet  10.1.13.2 <<< Incorrect IP address
                               iso
                               mpls
lo0               up    up
lo0.0             up    up    inet  10.0.0.1
                               iso  49.0004.1000.0000.0001.00

user@R3> show interfaces terse
Interface          Admin Link Proto Local Remote
so-0/0/0           up    up
so-0/0/0.0         up    up    inet  10.1.34.1/30
                               iso
                               mpls
so-0/0/1           up    up
so-0/0/1.0         up    up    inet  10.1.23.2/30
                               iso
                               mpls
so-0/0/2           up    up
so-0/0/2.0        up    up    inet  10.1.13.2/30 <<< Identical to R1
                               iso
                               mpls
so-0/0/3           up    up
so-0/0/3.0         up    up    inet  10.1.36.1/30
                               iso
                               mpls
lo0               up    up
lo0.0             up    up    inet  10.0.0.3
                               iso  49.0004.1000.0000.0003.00

user@R6> show interfaces terse
Interface          Admin Link Proto Local Remote
so-0/0/0           up    up
so-0/0/0.0         up    up    inet  10.1.56.2/30
                               iso
                               mpls
so-0/0/1           up    up
so-0/0/1.0         up    up    inet  10.1.46.2/30
                               iso

```



```

                                mpls
so-0/0/2                up    up
so-0/0/2.0              up    up    inet  10.1.26.2/30
                                iso
                                mpls
so-0/0/3                up    up
so-0/0/3.0              up    up    inet  10.1.36.2/30
                                iso
                                mpls
lo0.0                   up    up    inet  10.0.0.6
                                iso  49.0004.1000.0000.0006.00

```

Meaning

The sample output shows that the IP addresses for interface **so-0/0/2.0** on **R1** and interface **so-0/0/2.0** on **R3** are identical. Interface IP addresses within a network must be unique for the interface to be identified correctly.

Verify Neighbors or Adjacencies at the IP Layer

Purpose

If the IP addressing is configured incorrectly then the OSPF neighbors or IS-IS adjacencies both need to be checked to determine if one or both of them are established.

Action

To verify neighbors (OSPF) or adjacencies (IS-IS), enter the following commands from the ingress, transit, and egress routers:

```

user@host> show ospf neighbor extensive
user@host> show isis adjacency extensive

```

Sample Output 1

```

user@R1> show ospf neighbor extensive
Address          Interface          State      ID              Pri  Dead
10.1.12.2         so-0/0/0.0         Full      10.0.0.2        128  34
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 04:45:20, adjacent 1d 04:45:20
10.1.15.2         so-0/0/1.0         Full      10.0.0.5        128  35

```



```

area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
Up 1d 04:45:20, adjacent 1d 04:45:10 <<< no adjacency with R3 so-0/0/2

user@R3> show ospf neighbor extensive
Address          Interface          State      ID              Pri  Dead
10.1.23.1         so-0/0/1.0         Full       10.0.0.2        128  35
    area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
    Up 1w2d 04:54:30, adjacent 1w2d 04:54:21
10.1.36.2         so-0/0/3.0         Full       10.0.0.6        128  39
    area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
    Up 1w2d 04:54:30, adjacent 1w2d 04:54:30 <<< no adjacency with R1 so-0/0/2

user@R6> show ospf neighbor extensive
Address          Interface          State      ID              Pri  Dead
10.1.56.1         so-0/0/0.0         Full       10.0.0.5        128  39
    area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
    Up 1d 02:59:35, adjacent 1d 02:59:35
10.1.26.1         so-0/0/2.0         Full       10.0.0.2        128  36
    area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
    Up 1w2d 04:57:30, adjacent 1w2d 04:57:30
10.1.36.1         so-0/0/3.0         Full       10.0.0.3        128  36
    area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
    Up 1w2d 04:56:11, adjacent 1w2d 04:56:11

```

Sample Output 2

```

user@R1> show isis adjacency extensive
R2
    Interface: so-0/0/0.0,  Level: 2, State: Up , Expires in 23 secs
    Priority: 0, Up/Down transitions: 1, Last transition: 05:57:16 ago
    Circuit type: 2,  Speaks: IP , IPv6
    Topologies: Unicast
    Restart capable: Yes
    IP addresses: 10.1.12.2
    Transition log:
    When              State      Reason
    Fri Oct 15 14:58:35  Up        Seenself

R5

```


Interface: so-0/0/1.0, **Level: 2, State: Up**, Expires in 26 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 05:56:52 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.15.2

Transition log:

When	State	Reason
Fri Oct 15 14:59:00	Up	Seenself

R3

Interface: so-0/0/2.0, Level: 2, State: Up, Expires in 26 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 05:56:51 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.13.2

Transition log:

When	State	Reason
Fri Oct 15 14:59:01	Up	Seenself

user@R3> **show isis adjacency extensive**

R4

Interface: so-0/0/0.0, **Level: 2, State: Up** , Expires in 25 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:22:51 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.34.2

Transition log:

When	State	Reason
Thu Oct 28 15:13:12	Up	Seenself

R2

Interface: so-0/0/1.0, **Level: 2, State: Up** , Expires in 25 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 18:02:48 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.23.1

Transition log:

When	State	Reason
Tue Oct 19 21:33:15	Up	Seenself

R1

Interface: so-0/0/2.0, **Level: 2, State: Up** , Expires in 22 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 17:24:06 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.13.1

Transition log:

When	State	Reason
Tue Oct 19 22:11:57	Up	Seenself

R6

Interface: so-0/0/3.0, **Level: 2, State: Up** , Expires in 21 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:07:00 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.36.2

Transition log:

When	State	Reason
Thu Oct 21 15:29:03	Up	Seenself

user@R6> **show isis adjacency extensive**

R5

Interface: so-0/0/0.0, **Level: 2, State: Up** , Expires in 23 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 1w2d 01:10:03 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.56.1

Transition log:

When	State	Reason
Wed Oct 27 14:35:32	Up	Seenself

R4

Interface: so-0/0/1.0, **Level: 2, State: Up** , Expires in 25 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:26:50 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes
IP addresses: 10.1.46.1

Transition log:

When	State	Reason
Thu Oct 28 15:18:45	Up	Seenself

R2

Interface: so-0/0/2.0, **Level: 2, State: Up** , Expires in 24 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.26.1

Transition log:

When	State	Reason
Thu Oct 21 15:33:55	Up	Seenself

R3

Interface: so-0/0/3.0, **Level: 2, State: Up** , Expires in 19 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.36.1

Transition log:

When	State	Reason
Thu Oct 21 15:33:55	Up	Seenself

Meaning

Sample Output 1 from the ingress, transit, and egress routers shows that **R1** and **R3** are not established OSPF neighbors. Considering that the two interfaces **so-0/0/2.0** (**R1** and **R3**) are configured with identical IP addresses, you would expect this. The OSPF protocol routes IP packets based solely on the destination IP address contained in the IP packet header. Therefore, identical IP addresses in the autonomous system (AS) result in neighbors not establishing.

Sample Output 2 from the ingress, transit, and egress routers shows that **R1** and **R3** have established an IS-IS adjacency despite the identical IP addresses configured on interfaces **so-0/0/2.0** on **R1** and **R3**. The IS-IS protocol behaves differently from the OSPF protocol because it does not rely on IP to establish an adjacency. However, if the LSP is not up, it is still useful to check the IP subnet addressing in case there is a mistake in that layer. Correcting the addressing error might bring the LSP back up.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the IP address of an interface on transit router **R2** is incorrectly configured.

Solution

To correct the error in this example, enter the following commands:

```
[edit interfaces so-0/0/2]
user@R1# show
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30
user@R1# show
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.2/30; <<< Incorrect IP address
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30

[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.1/30; <<< Correct IP address.
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
user@R1# commit
commit complete
```


Meaning

The sample output shows that interface **so-0/0/2** on ingress router **R1** is now configured with the correct IP address. This correction results in unique subnet IP addresses for all interfaces in the MPLS network in [Figure 155 on page 1991](#), and the possibility that the LSP might come up.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the OSPF protocol has been resolved.

Action

To verify the LSP again, enter the following command on the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1 , LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
54 Oct 15 21:28:16 Selected as active path
53 Oct 15 21:28:16 Record Route: 10.1.13.2 10.1.36.2
52 Oct 15 21:28:16 Up
51 Oct 15 21:28:16 10.1.15.1: MPLS label allocation failure[2 times]
50 Oct 15 21:28:11 CSPF: computation result accepted
49 Oct 15 21:27:42 10.1.15.1: MPLS label allocation failure
48 Oct 15 21:27:42 CSPF: computation result accepted
47 Oct 15 21:27:31 10.1.15.1: MPLS label allocation failure[4 times]
```



```

    46 Oct 15 21:27:13 Originate Call
    45 Oct 15 21:27:13 CSPF: computation result accepted
    [...Output truncated...]
    Created: Thu Oct 14 16:04:34 2004
    Total 1 displayed,  Up 1 , Down 0

    Egress LSP: 1 sessions

10.0.0.1
    From: 10.0.0.6,  LSPstate: Up , ActiveRoute: 0
    LSPname: R6-to-R1 , LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: -
    Resv style: 1 FF, Label in: 3, Label out: -
    Time left: 149, Since: Fri Oct 15 21:28:13 2004
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 13 receiver 39024 protocol 0
    PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
    Adspec: received MTU 1500
    PATH sentto: localclient
    RESV rcvfrom: localclient
    Record route: 10.1.36.2 10.1.13.2 <self>
    Total 1 displayed,  Up 1 , Down 0

    Transit LSP: 0 sessions
    Total 0 displayed, Up 0, Down 0

```

Sample Output 2

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1

```



```

From: 10.0.0.6,  LSPstate: Up , ActiveRoute: 1
LSPname: R6-to-R1 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100336, Label out: 3
Time left: 156, Since: Fri Oct 15 21:15:47 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 13 receiver 39024 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
From: 10.0.0.1,  LSPstate: Up , ActiveRoute: 1
LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100352, Label out: 3
Time left: 159, Since: Fri Oct 15 21:15:50 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 5 receiver 47901 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 11 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2

Total 2 displayed,  Up 2 , Down 0

```

Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6,  State: Up , ActiveRoute: 1,  LSPname: R6-to-R1

```



```

ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.36.1 10.1.13.1
    187 Oct 15 21:20:05 Selected as active path
    186 Oct 15 21:20:05 Record Route: 10.1.36.1 10.1.13.1
    185 Oct 15 21:20:05 Up
    184 Oct 15 21:20:05 Clear Call
    183 Oct 15 21:20:05 CSPF: computation result accepted
    182 Oct 15 21:20:05 CSPF: link down/deleted
10.1.13.2(R3.00/10.0.0.3)->10.1.13.2(R1.00/10.0.0.1)
    [...Output truncated...]
    Created: Tue Aug 17 12:18:33 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.6
    From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
    LSPname: R1-to-R6 , LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: -
    Resv style: 1 FF, Label in: 3, Label out: -
    Time left: 144, Since: Fri Oct 15 21:20:08 2004
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 5 receiver 47901 protocol 0
    PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
    Adspec: received MTU 1500
    PATH sentto: localclient
    RESV rcvfrom: localclient
    Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1 , Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up. The output shows that the egress LSP session **R6-to-R1** received and sent a recovery label.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Verify the LSP

Purpose

After configuring the LSP, you must verify that the LSP is up. LSPs can be ingress, transit, or egress. Use the **show mpls lsp** command for quick verification of the LSP state, with the **extensive** option (**show mpls lsp extensive**) as a follow-up if the LSP is down. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To verify that the LSP is up, enter the following command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 25 second(s).
  44 Oct 15 16:56:11 CSPF failed:  no route toward 10.0.0.6 [2685 times]
  43 Oct 14 19:07:09 Clear Call
  42 Oct 14 19:06:56 Deselected as active
```



```

41 Oct 14 19:06:56 10.1.12.1: MPLS label allocation failure
40 Oct 14 19:06:56 Down
39 Oct 14 18:43:43 Selected as active path
38 Oct 14 18:43:43 Record Route: 10.1.13.2 10.1.36.2
37 Oct 14 18:43:43 Up
[...Output truncated...]
Created: Thu Oct 14 16:04:33 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed , Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed , Up 0, Down 0

```

Meaning

The sample output from ingress router **R1** shows that an MPLS label allocation failure occurred and the Constrained Shortest Path First (CSPF) algorithm failed, resulting in no route to destination **10.0.0.6** on **R6**.

Verify IP Addressing

Purpose

When you investigate the IP layer, you verify that interfaces have correct IP addressing, and that OSPF neighbors or IS-IS adjacencies are established. In this example, an IP address is configured incorrectly on the ingress router (**R1**).

Action

To verify IP addressing, enter the following command from the ingress, transit, and egress routers:

```
user@host> show interfaces terse
```

Sample Output

```

user@R1> show interfaces terse
Interface          Admin Link Proto Local          Remote

```



```

so-0/0/0          up    up
so-0/0/0.0        up    up    inet  10.1.12.1/30
                               iso
                               mpls

so-0/0/1          up    up
so-0/0/1.0        up    up    inet  10.1.15.1/30
                               iso
                               mpls

so-0/0/2          up    up
so-0/0/2.0        up    up    inet  10.1.13.2 <<< Incorrect IP address
                               iso
                               mpls

lo0               up    up
lo0.0             up    up    inet  10.0.0.1
                               iso  49.0004.1000.0000.0001.00

user@R3> show interfaces terse
Interface          Admin Link Proto Local              Remote
so-0/0/0           up    up
so-0/0/0.0         up    up    inet  10.1.34.1/30
                               iso
                               mpls

so-0/0/1           up    up
so-0/0/1.0         up    up    inet  10.1.23.2/30
                               iso
                               mpls

so-0/0/2           up    up
so-0/0/2.0         up    up    inet  10.1.13.2/30 <<< Identical to R1
                               iso
                               mpls

so-0/0/3           up    up
so-0/0/3.0         up    up    inet  10.1.36.1/30
                               iso
                               mpls

lo0               up    up
lo0.0             up    up    inet  10.0.0.3
                               iso  49.0004.1000.0000.0003.00

user@R6> show interfaces terse
Interface          Admin Link Proto Local              Remote
so-0/0/0           up    up
so-0/0/0.0         up    up    inet  10.1.56.2/30
                               iso
                               mpls

```



```

so-0/0/1          up    up
so-0/0/1.0        up    up    inet  10.1.46.2/30
                               iso
                               mpls

so-0/0/2          up    up
so-0/0/2.0        up    up    inet  10.1.26.2/30
                               iso
                               mpls

so-0/0/3          up    up
so-0/0/3.0        up    up    inet  10.1.36.2/30
                               iso
                               mpls

lo0.0             up    up    inet  10.0.0.6
                               iso  49.0004.1000.0000.0006.00

```

Meaning

The sample output shows that the IP addresses for interface **so-0/0/2.0** on **R1** and interface **so-0/0/2.0** on **R3** are identical. Interface IP addresses within a network must be unique for the interface to be identified correctly.

Verify Neighbors or Adjacencies at the IP Layer

Purpose

If the IP addressing is configured incorrectly then the OSPF neighbors or IS-IS adjacencies both need to be checked to determine if one or both of them are established.

Action

To verify neighbors (OSPF) or adjacencies (IS-IS), enter the following commands from the ingress, transit, and egress routers:

```

user@host> show ospf neighbor extensive
user@host> show isis adjacency extensive

```

Sample Output 1

```

user@R1> show ospf neighbor extensive

```

Address	Interface	State	ID	Pri	Dead
---------	-----------	-------	----	-----	------


```

10.1.12.2      so-0/0/0.0      Full    10.0.0.2      128    34
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 04:45:20, adjacent 1d 04:45:20
10.1.15.2      so-0/0/1.0      Full    10.0.0.5      128    35
  area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0
  Up 1d 04:45:20, adjacent 1d 04:45:10 <<< no adjacency with R3 so-0/0/2

```

user@R3> show ospf neighbor extensive

Address	Interface	State	ID	Pri	Dead
10.1.23.1	so-0/0/1.0	Full	10.0.0.2	128	35
area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0					
Up 1w2d 04:54:30, adjacent 1w2d 04:54:21					
10.1.36.2	so-0/0/3.0	Full	10.0.0.6	128	39
area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0					
Up 1w2d 04:54:30, adjacent 1w2d 04:54:30 <<< no adjacency with R1 so-0/0/2					

user@R6> show ospf neighbor extensive

Address	Interface	State	ID	Pri	Dead
10.1.56.1	so-0/0/0.0	Full	10.0.0.5	128	39
area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0					
Up 1d 02:59:35, adjacent 1d 02:59:35					
10.1.26.1	so-0/0/2.0	Full	10.0.0.2	128	36
area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0					
Up 1w2d 04:57:30, adjacent 1w2d 04:57:30					
10.1.36.1	so-0/0/3.0	Full	10.0.0.3	128	36
area 0.0.0.0, opt 0x42, DR 0.0.0.0, BDR 0.0.0.0					
Up 1w2d 04:56:11, adjacent 1w2d 04:56:11					

Sample Output 2

user@R1> show isis adjacency extensive

R2

```

Interface: so-0/0/0.0,  Level: 2, State: Up , Expires in 23 secs
Priority: 0, Up/Down transitions: 1, Last transition: 05:57:16 ago
Circuit type: 2,  Speaks: IP , IPv6
Topologies: Unicast
Restart capable: Yes
  IP addresses: 10.1.12.2
Transition log:

```


When	State	Reason
Fri Oct 15 14:58:35	Up	Seenself

R5

Interface: so-0/0/1.0, **Level: 2, State: Up**, Expires in 26 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 05:56:52 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.15.2

Transition log:

When	State	Reason
Fri Oct 15 14:59:00	Up	Seenself

R3

Interface: so-0/0/2.0, Level: 2, State: Up, Expires in 26 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 05:56:51 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.13.2

Transition log:

When	State	Reason
Fri Oct 15 14:59:01	Up	Seenself

user@R3> show isis adjacency extensive

R4

Interface: so-0/0/0.0, **Level: 2, State: Up** , Expires in 25 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:22:51 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.34.2

Transition log:

When	State	Reason
Thu Oct 28 15:13:12	Up	Seenself

R2

Interface: so-0/0/1.0, **Level: 2, State: Up** , Expires in 25 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 18:02:48 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.23.1

Transition log:

When	State	Reason
Tue Oct 19 21:33:15	Up	Seenself

R1

Interface: so-0/0/2.0, **Level: 2, State: Up** , Expires in 22 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w2d 17:24:06 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.13.1

Transition log:

When	State	Reason
Tue Oct 19 22:11:57	Up	Seenself

R6

Interface: so-0/0/3.0, **Level: 2, State: Up** , Expires in 21 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:07:00 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.36.2

Transition log:

When	State	Reason
Thu Oct 21 15:29:03	Up	Seenself

user@R6> **show isis adjacency extensive**

R5

Interface: so-0/0/0.0, **Level: 2, State: Up** , Expires in 23 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 1w2d 01:10:03 ago
 Circuit type: 2, **Speaks: IP** , IPv6
 Topologies: Unicast
 Restart capable: Yes

IP addresses: 10.1.56.1

Transition log:

When	State	Reason
Wed Oct 27 14:35:32	Up	Seenself

R4

Interface: so-0/0/1.0, **Level: 2, State: Up** , Expires in 25 secs
 Priority: 0, Up/Down transitions: 1, Last transition: 1w1d 00:26:50 ago
 Circuit type: 2, **Speaks: IP** , IPv6

Topologies: Unicast
Restart capable: Yes

IP addresses: 10.1.46.1

Transition log:

When	State	Reason
Thu Oct 28 15:18:45	Up	Seenself

R2

Interface: so-0/0/2.0, **Level: 2, State: Up** , Expires in 24 secs
Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
Circuit type: 2, **Speaks: IP** , IPv6
Topologies: Unicast
Restart capable: Yes

IP addresses: 10.1.26.1

Transition log:

When	State	Reason
Thu Oct 21 15:33:55	Up	Seenself

R3

Interface: so-0/0/3.0, **Level: 2, State: Up** , Expires in 19 secs
Priority: 0, Up/Down transitions: 1, Last transition: 2w1d 00:11:40 ago
Circuit type: 2, **Speaks: IP** , IPv6
Topologies: Unicast
Restart capable: Yes

IP addresses: 10.1.36.1

Transition log:

When	State	Reason
Thu Oct 21 15:33:55	Up	Seenself

Meaning

Sample Output 1 from the ingress, transit, and egress routers shows that **R1** and **R3** are not established OSPF neighbors. Considering that the two interfaces **so-0/0/2.0** (**R1** and **R3**) are configured with identical IP addresses, you would expect this. The OSPF protocol routes IP packets based solely on the destination IP address contained in the IP packet header. Therefore, identical IP addresses in the autonomous system (AS) result in neighbors not establishing.

Sample Output 2 from the ingress, transit, and egress routers shows that **R1** and **R3** have established an IS-IS adjacency despite the identical IP addresses configured on interfaces **so-0/0/2.0** on **R1** and **R3**. The IS-IS protocol behaves differently from the OSPF protocol because it does not rely on IP to establish an adjacency. However, if the LSP is not up, it is still useful to check the IP subnet addressing in case there is a mistake in that layer. Correcting the addressing error might bring the LSP back up.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the IP address of an interface on transit router **R2** is incorrectly configured.

Solution

To correct the error in this example, enter the following commands:

```
[edit interfaces so-0/0/2]
user@R1# show
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30
user@R1# show
user@R1# commit
```

Sample Output

```
[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.2/30; <<< Incorrect IP address
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
user@R1# rename unit 0 family inet address 10.1.13.2/30 to address 10.1.13.1/30

[edit interfaces so-0/0/2]
user@R1# show
unit 0 {
    family inet {
        address 10.1.13.1/30; <<< Correct IP address.
    }
    family iso;
    family mpls;
}

[edit interfaces so-0/0/2]
```



```
user@R1# commit
commit complete
```

Meaning

The sample output shows that interface **so-0/0/2** on ingress router **R1** is now configured with the correct IP address. This correction results in unique subnet IP addresses for all interfaces in the MPLS network in [Figure 155 on page 1991](#), and the possibility that the LSP might come up.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the OSPF protocol has been resolved.

Action

To verify the LSP again, enter the following command on the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1 , LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    54 Oct 15 21:28:16 Selected as active path
```



```

53 Oct 15 21:28:16 Record Route: 10.1.13.2 10.1.36.2
52 Oct 15 21:28:16 Up
51 Oct 15 21:28:16 10.1.15.1: MPLS label allocation failure[2 times]
50 Oct 15 21:28:11 CSPF: computation result accepted
49 Oct 15 21:27:42 10.1.15.1: MPLS label allocation failure
48 Oct 15 21:27:42 CSPF: computation result accepted
47 Oct 15 21:27:31 10.1.15.1: MPLS label allocation failure[4 times]
46 Oct 15 21:27:13 Originate Call
45 Oct 15 21:27:13 CSPF: computation result accepted
[...Output truncated...]
Created: Thu Oct 14 16:04:34 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up , ActiveRoute: 0
  LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 149, Since: Fri Oct 15 21:28:13 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 13 receiver 39024 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1 , Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 2

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```


Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1

From: 10.0.0.6, **LSPstate: Up** , ActiveRoute: 1
LSPname: R6-to-R1 , LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100336, Label out: 3
 Time left: 156, Since: Fri Oct 15 21:15:47 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 13 receiver 39024 protocol 0
 PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
 RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
 Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6

From: 10.0.0.1, **LSPstate: Up** , ActiveRoute: 1
LSPname: R1-to-R6 , LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100352, Label out: 3
 Time left: 159, Since: Fri Oct 15 21:15:50 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 5 receiver 47901 protocol 0
 PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.36.2 (so-0/0/3.0) 11 pkts
 RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
 Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2

Total 2 displayed, **Up 2** , Down 0

Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Up , ActiveRoute: 1, LSPname: R6-to-R1
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.36.1 10.1.13.1
    187 Oct 15 21:20:05 Selected as active path
    186 Oct 15 21:20:05 Record Route:  10.1.36.1 10.1.13.1
    185 Oct 15 21:20:05 Up
    184 Oct 15 21:20:05 Clear Call
    183 Oct 15 21:20:05 CSPF: computation result accepted
    182 Oct 15 21:20:05 CSPF: link down/deleted
10.1.13.2(R3.00/10.0.0.3)->10.1.13.2(R1.00/10.0.0.1)
  [...Output truncated...]
  Created: Tue Aug 17 12:18:33 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 144, Since: Fri Oct 15 21:20:08 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 47901 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>

```



```
Total 1 displayed,  Up 1 , Down 0
```

```
Transit LSP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

Meaning

Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up. The output shows that the egress LSP session **R6-to-R1** received and sent a recovery label.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

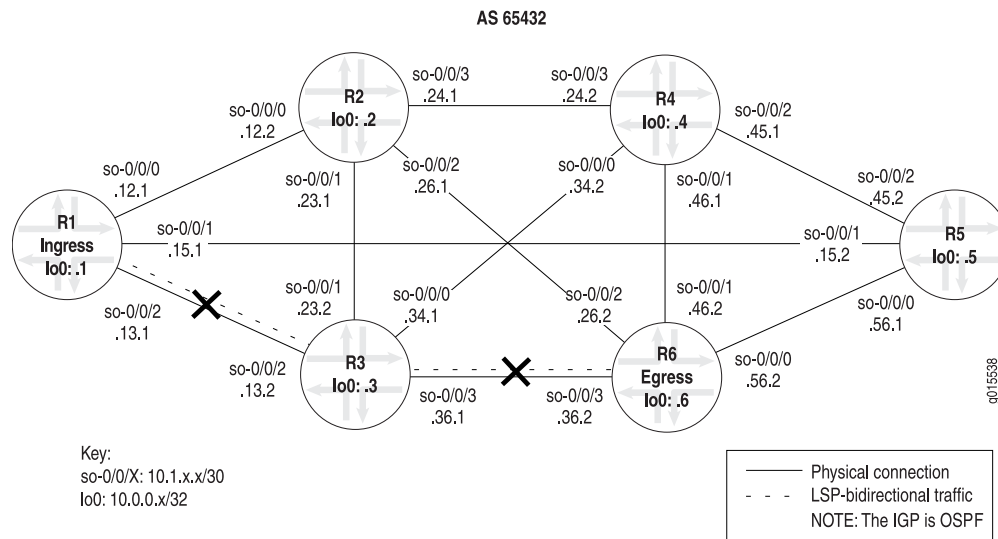
Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Verifying the OSPF Protocol

Purpose

After you have verified that the LSP is down, and the cause is not in the physical, datalink, or IP layer, verify the OSPF configuration. Check the routers in your network to ensure that the interfaces and the OSPF protocol are configured correctly, and that the neighbors are established.

Figure 157: MPLS Network Broken at the OSPF Protocol Layer



1. [Verify the LSP | 2020](#)
2. [Verify OSPF Interfaces | 2025](#)
3. [Verify OSPF Neighbors | 2027](#)
4. [Verify the OSPF Protocol Configuration | 2028](#)
5. [Take Appropriate Action | 2029](#)
6. [Verify the LSP Again | 2031](#)

Verify the LSP

Purpose

Confirm that interfaces are configured for OSPF, the OSPF protocol is configured correctly and that neighbors are established.

Action

To verify the LSP, enter the following command on the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```


Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0,  LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    11 Oct 19 18:06:04 No Route toward dest[78 times]
    10 Oct 19 17:08:09 Deselected as active
  Created: Mon Oct 18 21:48:42 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```


Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.1     10.0.0.6      Dn     0  -              R6-to-R1
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 4

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1,  LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    5 Oct 19 10:37:55 Selected as active path
    4 Oct 19 10:37:55 Record Route:  10.1.13.2 10.1.36.2
    3 Oct 19 10:37:55 Up
    2 Oct 19 10:37:10 No Route toward dest[1029 times]
    1 Oct 18 21:48:42 Originate Call
    Created: Mon Oct 18 21:48:42 2004
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
```

Sample Output 5

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
    LSPname: R1-to-R6 , LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: 3
    Resv style: 1 FF, Label in: 100368, Label out: 3
    Time left: 154, Since: Tue Oct 19 10:25:24 2004
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 1 receiver 47933 protocol 0
    PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 209 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.1.36.2 (so-0/0/3.0) 209 pkts
    RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 209 pkts
    Record route: 10.1.13.1 <self> 10.1.36.2
Total 1 displayed, Up 1, Down 0
```

Sample Output 6

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
```



```

From: 10.0.0.6, State: Dn, ActiveRoute: 0,  LSPname: R6-to-R1
  ActivePath: (none)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary                               State: Dn
  2 Oct 19 13:01:54 10.1.56.2:  MPLS label allocation failure [9 times]
  1 Oct 19 12:57:51 Originate Call
Created: Tue Oct 19 12:57:51 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 148, Since: Tue Oct 19 10:30:03 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47933 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 206 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Outputs 1, 2, and 3 show that the LSP and the reverse LSP are down:

- Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** does not have a route towards the destination (**R6**).
- Sample Output 2 from transit router **R3** shows that there are no LSP sessions.
- Sample Output 3 from egress router **R6** also shows that reverse LSP **R6-to-R1** is down.

Sample Outputs 4, 5, and 6 show that the LSP is up and the reverse LSP is down:

- Sample Output 4 from ingress router **R1** shows that LSP **R1-to-R6** is up and there are no egress LSP sessions.

- Sample Output 5 from transit router R3 shows that there is one ingress LSP session (**R1-to-R6**) and no egress LSP sessions.
- Sample Output 6 from egress router **R6** shows that LSP **R6-to-R1** is down due to an MPLS label allocation failure.

Verify OSPF Interfaces

Purpose

After you have verified that the LSP is down, and the cause is not in the physical, data link, or IP layer, check the routers in your network to determine that all relevant OSPF interfaces are configured correctly.

Action

To verify OSPF interfaces, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf interface
```

Sample Output 1

```
user@R1> show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

```
user@R3> show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

```
user@R6> show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Sample Output 2

```
user@R1> show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.0.0.1	0.0.0.0	0
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

```
user@R3> show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.0.0.3	0.0.0.0	0
so-0/0/0.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

```
user@R6> show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.0.0.6	0.0.0.0	0
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Meaning

Sample Output 1 shows that all interfaces on all routers are in the correct area (0.0.0.0), and the loopback (lo0) interface is missing from the list of interfaces on all routers. The missing loopback (lo0) interface is a problem in this configuration.

In an MPLS network configured with OSPF as the IGP, when you manually configure the RID, it is important to explicitly configure the loopback interface at the `[edit protocols ospf]` hierarchy level. If the RID is not manually configured, OSPF automatically advertises the loopback (lo0) interface. In the configuration of all the routers in this network, the RID is configured manually, therefore, the loopback (lo0) interface must be explicitly configured at the `[edit protocols ospf]` hierarchy level. In addition, the loopback (lo0) interface is configured with the **passive** statement to ensure that the protocols are not run over the loopback (lo0) interface and it is correctly advertised throughout the network.

Sample Output 2 shows that all the relevant interfaces on the ingress, transit, and egress routers, including the loopback (lo0) interface, are in the correct area (0.0.0.0). Because the configuration of the interfaces is correct, further investigation is required to determine the reason for the LSP problem.

Verify OSPF Neighbors

Purpose

After you have checked OSPF interfaces, check your network topology to determine that all relevant neighbors are established.

Action

To verify OSPF neighbors, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf neighbor
```

Sample Output

```
user@R1> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.12.2	so-0/0/0.0	Full	10.0.0.2	128	39
10.1.15.2	so-0/0/1.0	Full	10.0.0.5	128	39
10.1.13.2	so-0/0/2.0	Full	10.0.0.3	128	33

```
user@R3> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.34.2	so-0/0/0.0	Full	10.0.0.4	128	33
10.1.23.1	so-0/0/1.0	Full	10.0.0.2	128	33
10.1.13.1	so-0/0/2.0	Full	10.0.0.1	128	33
10.1.36.2	so-0/0/3.0	Full	10.0.0.6	128	33

```
user@R6> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.56.1	so-0/0/0.0	Full	10.0.0.5	128	30
10.1.46.1	so-0/0/1.0	Full	10.0.0.4	128	38
10.1.26.1	so-0/0/2.0	Full	10.0.0.2	128	34
10.1.36.1	so-0/0/3.0	Full	10.0.0.3	128	35

Meaning

The sample output shows that all neighbors are fully adjacent, indicating that each router has exchanged a full copy of its link-state database with the other routers, passed through several neighbor states, and become fully adjacent. These adjacencies are created by router link and network link advertisements.

Verify the OSPF Protocol Configuration

Purpose

After you have checked interfaces and neighbors, verify the OSPF protocol configuration.

Action

To verify the OSPF protocol configuration, enter the following command from the ingress, transit, and egress routers:

```
user@host> show configuration protocols ospf
```

Sample Output 1

```
user@R1> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;    <<< The loopback interface (lo0) is missing
}
```

Sample Output 2

```
user@R3>show configuration protocols ospf
area 0.0.0.0 { <<< traffic engineering is missing
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface so-0/0/3.0;    <<< The loopback interface (lo0) is missing
}
```


Sample Output 3

```

user@R6> show configuration protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface so-0/0/3.0;    <<< The loopback interface (lo0) is missing
}

```

Meaning

All three sample outputs show that the loopback interface is not included on any of the routers. Including the loopback (**lo0**) interface is important when you have the RID manually configured.

In addition, Sample Output 2 from transit router **R3** shows that traffic engineering is not configured. Traffic engineering must be manually enabled when you configure OSPF for an MPLS network.

Because the loopback interface and traffic engineering are missing from the OSPF protocol configuration, the LSP does not work as expected.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, the loopback (**lo0**) interface is missing from all routers, and traffic engineering is missing from the transit router (**R3**).

Solution

To correct the errors in this example, follow these steps:

1. Include the loopback (**lo0**) interface on all routers that have the RID manually configured. Enter the following configuration mode commands:

```

[edit]
user@R3# edit protocols ospf area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0
user@R3# set interface lo0 passive

```

2. Move up one level of the configuration hierarchy:


```
[edit protocols ospf area 0.0.0.0]
user@R3# up
[edit protocols ospf]
user @R3#
```

3. Include traffic engineering on the transit router (**R3**). Enter the following configuration mode command:

```
[edit protocols ospf]
user@R3# set traffic-engineering
```

4. On all routers, verify and commit the configuration:

```
user@R3# show
user@R3# commit
```

Sample Output

```
user@R3> edit
Entering configuration mode

[edit]
user@R3# edit protocols ospf area 0.0.0.0

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0

[edit protocols ospf area 0.0.0.0]
user@R3# set interface lo0 passive

[edit protocols ospf area 0.0.0.0]
user@R3# up

[edit protocols ospf]
user@R3# set traffic-engineering

[edit protocols ospf]
user@R3# show
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
```



```

interface so-0/0/2.0;
  interface lo0.0; {
    passive
  }
}

[edit protocols ospf]
user@R3# commit
commit complete

```

Meaning

The sample output shows that the loopback (**lo0**) interface and traffic engineering are now correctly configured on transit router **R3**. When traffic engineering is configured, OSPF advertises the traffic engineering capabilities of the links.

In the OSPF configuration, you must manually include the loopback (**lo0**) interface and set it to passive when you manually configure an RID. Setting the loopback (**lo0**) interface to passive ensures that protocols are not run over the loopback (**lo0**) interface and the loopback (**lo0**) interface is advertised correctly throughout the network.. If you do not manually configure an RID, there is no need to explicitly include the loopback interface because the OSPF protocol automatically includes the loopback (**lo0**) interface.

For more information about configuring LSPs and MPLS, see the *Junos MPLS Applications Configuration Guide*.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the IS-IS protocol has been resolved.

Action

To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

```
user@host> show mpls lsp extensive
```


Sample Output

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up , ActiveRoute: 1, LSPname: R1-to-R6
    ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    4 Oct 19 21:22:54 Selected as active path
    3 Oct 19 21:22:53 Record Route:  10.1.13.2 10.1.36.2
    2 Oct 19 21:22:53 Up
    1 Oct 19 21:22:53 Originate Call
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up , ActiveRoute: 0
    LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 117, Since: Tue Oct 19 21:17:42 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39064 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1 , Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive

```



```
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Transit LSP: 2 sessions
```

```
10.0.0.1
```

```
From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100416, Label out: 3
Time left: 139, Since: Tue Oct 19 21:05:11 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 39064 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1
```

```
10.0.0.6
```

```
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100448, Label out: 3
Time left: 135, Since: Tue Oct 19 21:10:22 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47951 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
Record route: 10.1.13.1 <self> 10.1.36.2
```

```
Total 2 displayed, Up 2 , Down 0
```

```
user@R6> run show mpls lsp extensive
```

```
Ingress LSP: 1 sessions
```

```
10.0.0.1
```



```

From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.36.1 10.1.13.1
19 Oct 19 21:09:52 Selected as active path
18 Oct 19 21:09:52 Record Route: 10.1.36.1 10.1.13.1
17 Oct 19 21:09:52 Up
16 Oct 19 21:09:52 Originate Call
15 Oct 19 21:09:52 CSPF: computation result accepted
Created: Tue Oct 19 18:30:09 2004
Total 1 displayed,  Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.6
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 120, Since: Tue Oct 19 21:15:03 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47951 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed,  Up 1 , Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output from ingress router **R1** and egress router **R6** shows that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to

R1. In addition, the sample output from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6**, and the other from **R6** to **R1**.

Verify the LSP

Purpose

Confirm that interfaces are configured for OSPF, the OSPF protocol is configured correctly and that neighbors are established.

Action

To verify the LSP, enter the following command on the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0,  LSPname: R1-to-R6
ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    11 Oct 19 18:06:04 No Route toward dest[78 times]
    10 Oct 19 17:08:09 Deselected as active
  Created: Mon Oct 18 21:48:42 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```


Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.1     10.0.0.6      Dn     0  -                R6-to-R1
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 4

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1,  LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
```



```

*Primary                               State: Up
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
5 Oct 19 10:37:55 Selected as active path
4 Oct 19 10:37:55 Record Route:  10.1.13.2 10.1.36.2
3 Oct 19 10:37:55 Up
2 Oct 19 10:37:10 No Route toward dest[1029 times]
1 Oct 18 21:48:42 Originate Call
Created: Mon Oct 18 21:48:42 2004
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

```

Sample Output 5

```

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 1 sessions

10.0.0.6
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100368, Label out: 3
Time left:  154, Since: Tue Oct 19 10:25:24 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47933 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 209 pkts
Adspec: received MTU 1500 sent MTU 1500

```



```

PATH sentto: 10.1.36.2 (so-0/0/3.0) 209 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 209 pkts
Record route: 10.1.13.1 <self> 10.1.36.2
Total 1 displayed, Up 1, Down 0

```

Sample Output 6

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0,  LSPname: R6-to-R1
    ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    2 Oct 19 13:01:54 10.1.56.2:  MPLS label allocation failure [9 times]
    1 Oct 19 12:57:51 Originate Call
  Created: Tue Oct 19 12:57:51 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 148, Since: Tue Oct 19 10:30:03 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47933 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 206 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

```



```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

Sample Outputs 1, 2, and 3 show that the LSP and the reverse LSP are down:

- Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** does not have a route towards the destination (**R6**).
- Sample Output 2 from transit router **R3** shows that there are no LSP sessions.
- Sample Output 3 from egress router **R6** also shows that reverse LSP **R6-to-R1** is down.

Sample Outputs 4, 5, and 6 show that the LSP is up and the reverse LSP is down:

- Sample Output 4 from ingress router **R1** shows that LSP **R1-to-R6** is up and there are no egress LSP sessions.
- Sample Output 5 from transit router **R3** shows that there is one ingress LSP session (**R1-to-R6**) and no egress LSP sessions.
- Sample Output 6 from egress router **R6** shows that LSP **R6-to-R1** is down due to an MPLS label allocation failure.

Verify OSPF Interfaces

Purpose

After you have verified that the LSP is down, and the cause is not in the physical, data link, or IP layer, check the routers in your network to determine that all relevant OSPF interfaces are configured correctly.

Action

To verify OSPF interfaces, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf interface
```

Sample Output 1

```
user@R1> show ospf interface
```

Interface	State	Area	DR ID	BDR ID	Nbrs
-----------	-------	------	-------	--------	------


```
so-0/0/0.0    PtToPt  0.0.0.0    0.0.0.0    0.0.0.0    1
so-0/0/1.0    PtToPt  0.0.0.0    0.0.0.0    0.0.0.0    1
so-0/0/2.0    PtToPt  0.0.0.0    0.0.0.0    0.0.0.0    1
```

user@R3> show ospf interface

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

user@R6> show ospf interface

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Sample Output 2

user@R1> show ospf interface

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.0.0.1	0.0.0.0	0
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

user@R3> show ospf interface

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.0.0.3	0.0.0.0	0
so-0/0/0.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-0/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

user@R6> show ospf interface

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.0.0.6	0.0.0.0	0
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

so-0/0/1.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-0/0/2.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Meaning

Sample Output 1 shows that all interfaces on all routers are in the correct area (**0.0.0.0**), and the loopback (**lo0**) interface is missing from the list of interfaces on all routers. The missing loopback (**lo0**) interface is a problem in this configuration.

In an MPLS network configured with OSPF as the IGP, when you manually configure the RID, it is important to explicitly configure the loopback interface at the **[edit protocols ospf]** hierarchy level. If the RID is not manually configured, OSPF automatically advertises the loopback (**lo0**) interface. In the configuration of all the routers in this network, the RID is configured manually, therefore, the loopback (**lo0**) interface must be explicitly configured at the **[edit protocols ospf]** hierarchy level. In addition, the loopback (**lo0**) interface is configured with the **passive** statement to ensure that the protocols are not run over the loopback (**lo0**) interface and it is correctly advertised throughout the network.

Sample Output 2 shows that all the relevant interfaces on the ingress, transit, and egress routers, including the loopback (**lo0**) interface, are in the correct area (**0.0.0.0**). Because the configuration of the interfaces is correct, further investigation is required to determine the reason for the LSP problem.

Verify OSPF Neighbors

Purpose

After you have checked OSPF interfaces, check your network topology to determine that all relevant neighbors are established.

Action

To verify OSPF neighbors, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show ospf neighbor
```

Sample Output

```
user@R1> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.12.2	so-0/0/0.0	Full	10.0.0.2	128	39


```

10.1.15.2      so-0/0/1.0      Full    10.0.0.5      128    39
10.1.13.2      so-0/0/2.0      Full    10.0.0.3      128    33

```

```
user@R3> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.34.2	so-0/0/0.0	Full	10.0.0.4	128	33
10.1.23.1	so-0/0/1.0	Full	10.0.0.2	128	33
10.1.13.1	so-0/0/2.0	Full	10.0.0.1	128	33
10.1.36.2	so-0/0/3.0	Full	10.0.0.6	128	33

```
user@R6> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.56.1	so-0/0/0.0	Full	10.0.0.5	128	30
10.1.46.1	so-0/0/1.0	Full	10.0.0.4	128	38
10.1.26.1	so-0/0/2.0	Full	10.0.0.2	128	34
10.1.36.1	so-0/0/3.0	Full	10.0.0.3	128	35

Meaning

The sample output shows that all neighbors are fully adjacent, indicating that each router has exchanged a full copy of its link-state database with the other routers, passed through several neighbor states, and become fully adjacent. These adjacencies are created by router link and network link advertisements.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the IS-IS protocol has been resolved.

Action

To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

```
user@host> show mpls lsp extensive
```


Sample Output

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up , ActiveRoute: 1, LSPname: R1-to-R6
    ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    4 Oct 19 21:22:54 Selected as active path
    3 Oct 19 21:22:53 Record Route:  10.1.13.2 10.1.36.2
    2 Oct 19 21:22:53 Up
    1 Oct 19 21:22:53 Originate Call
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up , ActiveRoute: 0
    LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 117, Since: Tue Oct 19 21:17:42 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39064 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1 , Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive

```



```
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Transit LSP: 2 sessions
```

```
10.0.0.1
```

```
From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100416, Label out: 3
Time left: 139, Since: Tue Oct 19 21:05:11 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 39064 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 11 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 11 pkts
RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 11 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1
```

```
10.0.0.6
```

```
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100448, Label out: 3
Time left: 135, Since: Tue Oct 19 21:10:22 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47951 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
Record route: 10.1.13.1 <self> 10.1.36.2
```

```
Total 2 displayed, Up 2 , Down 0
```

```
user@R6> run show mpls lsp extensive
```

```
Ingress LSP: 1 sessions
```

```
10.0.0.1
```



```

From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.36.1 10.1.13.1
19 Oct 19 21:09:52 Selected as active path
18 Oct 19 21:09:52 Record Route: 10.1.36.1 10.1.13.1
17 Oct 19 21:09:52 Up
16 Oct 19 21:09:52 Originate Call
15 Oct 19 21:09:52 CSPF: computation result accepted
Created: Tue Oct 19 18:30:09 2004
Total 1 displayed,  Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.6
From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6 , LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 120, Since: Tue Oct 19 21:15:03 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47951 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed,  Up 1 , Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output from ingress router **R1** and egress router **R6** shows that the LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to

R1. In addition, the sample output from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6**, and the other from **R6** to **R1**.

Verify the LSP

Purpose

Confirm that interfaces are configured for IS-IS, that the IS-IS protocol is configured correctly, and that adjacencies are established.

Action

To verify the label-switched path (LSP), enter the following command on the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn,  ActiveRoute: 0 , LSPname: R1-to-R6
    ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    24 Oct 21 13:48:01  No Route toward dest [3 times]
    23 Oct 21 13:47:44 Deselected as active
    22 Oct 21 13:47:43 No Route toward dest[2 times]
    21 Oct 21 13:47:43  ResvTear received
    20 Oct 21 13:47:43 Down
    19 Oct 21 13:47:43 10.1.13.2: No Route toward dest[2 times]
    18 Oct 21 13:47:38 Record Route:  10.1.13.2 10.1.36.2
    [...Output truncated...]
  Created: Tue Oct 19 21:22:53 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
```



```
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 3

```
user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0 , LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 3 second(s).
    13 Oct 21 14:23:33 CSPF failed: no route toward 10.0.0.1[90 times]
    12 Oct 21 13:39:56 Deselected as active
    11 Oct 21 13:39:56 CSPF: could not determine self
    [...Output truncated...]
  Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 0, Down 1
```



```
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

The sample output shows that LSP **R1-to-R6** and the reverse LSP **R6-to-R1** are down, and there are no LSP sessions on transit router R3.

Verify IS-IS Adjacencies and Interfaces

Purpose

When you check the IS-IS layer, you verify that IS-IS adjacencies are up and that the IS-IS interfaces are included at the protocol level.

Action

To verify the functioning of adjacent interfaces, enter the following commands from the relevant routers:

```
user@host> show isis adjacency
user@host> show isis interface
```

Sample Output 1

```
user@R1> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R2	2 Up	20	
so-0/0/1.0	R5	2 Up	23	
so-0/0/2.0	R3	2 Up	26	

```
user@R3> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R4	2 Up	23	
so-0/0/1.0	R2	2 Up	21	
so-0/0/2.0	R1	2 Up	19	
so-0/0/3.0	R6	2 Down		0


```
user@R6> show isis adjacency
IS-IS instance is not running
```

Sample Output 2

```
user@R1> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0  0x1 Passive            Passive            0/0
so-0/0/0.0         2  0x1 Disabled          Point to Point    10/10
so-0/0/1.0         2  0x1 Disabled          Point to Point    10/10
so-0/0/2.0         2  0x1 Disabled          Point to Point    10/10

user@R3> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0  0x1 Passive            Passive            0/0
so-0/0/0.0         2  0x1 Disabled          Point to Point    10/10
so-0/0/1.0         2  0x1 Disabled          Point to Point    10/10
so-0/0/2.0         2  0x1 Disabled          Point to Point    10/10
so-0/0/3.0         2  0x1 Disabled          Point to Point    10/10

user@R6> show isis interface
IS-IS interface database:
Interface          L  CirID Level 1 DR      Level 2 DR      L1/L2 Metric
lo0.0              0  0x1 Passive            Passive            0/0
so-0/0/0.0         1   0x1 Point to Point    Disabled          10/10
so-0/0/1.0         1   0x1 Down              Disabled          10/10
so-0/0/2.0         1   0x1 Point to Point    Disabled          10/10
so-0/0/3.0         1   0x1 Point to Point    Disabled          10/10
```

Meaning

Sample Output 1 shows that ingress router R1 has established adjacencies with the relevant routers. Transit router R3 does not have an adjacency with egress router R6, and egress router R6 has no adjacencies established in the network shown in [Figure 155 on page 1991](#), indicating that the problem might be at the IS-IS protocol level.

Sample Output 2 shows that R1 and R2 are Level 2 routers, in contrast to R6 which is a Level 1 router. When a router is configured explicitly as a Level 1 or Level 2 router, it does not communicate with routers

configured at a different level. Level 1 routers communicate with other Level 1 routers within their area, while Level 2 routers communicate with other Level 2 routers, and toward other autonomous systems. Because all the routers in this network are configured for Level 2, they cannot form an adjacency with R6, which is incorrectly configured as a Level 1 router.

RELATED DOCUMENTATION

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

Verify the IS-IS Configuration

Purpose

When you have determined that the problem is probably at the IS-IS protocol level, check the IS-IS configuration of the routers in your network.

Action

To verify the IS-IS configuration, enter the following command from the relevant routers:

```
user@host> show configuration protocols isis
```

Sample Output

```
user@R1> show configuration protocols isis
```

```
level 1 disable;
```

```
interface so-0/0/0.0;
```

```
interface so-0/0/1.0;
```

```
interface so-0/0/2.0;
```

```
interface lo0.0; {  
    passive
```

```
user@R3> show configuration protocols isis
```

```
level 1 disable;
```

```
interface all {  
    level 2 metric 10;  
}
```



```

interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

user@R6> show configuration protocols isis
level 2 disable; <<< Incorrect level disabled
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

```

Meaning

The sample output shows that R6 has Level 2 disabled, while R1 and R3 have Level 1 disabled. For IS-IS adjacencies to establish, routers need to be at the same level. Another common configuration error is to omit the loopback interface (lo0) from the configuration at the **[edit protocols isis]** hierarchy level. IS-IS does not function correctly if the loopback interface (lo0) is not configured at this level. In addition, including the **passive** statement ensures that protocols are not run over the loopback interface (lo0) and that the loopback interface (lo0) is advertised correctly throughout the network.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the label-switched path (LSP) needs to be checked again to confirm that the problem in the RSVP layer has been resolved.

Action

To verify that the LSP is up and traversing the network as expected, enter the following command from the ingress, egress, and transit routers:

```
user@host> show mpls lsp extensive
```


Sample Output 1

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
    5 Oct 21 15:52:07 Selected as active path
    4 Oct 21 15:52:07 Record Route: 10.1.13.2 10.1.36.2
    3 Oct 21 15:52:07 Up
    2 Oct 21 15:52:07 Originate Call
    1 Oct 21 15:52:07 CSPF: computation result accepted
    Created: Thu Oct 21 15:52:06 2004
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 142, Since: Thu Oct 21 15:41:59 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 39082 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 17 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1 , Down 0

```



```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
    LSPname: R6-to-R1 , LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: 3
    Resv style: 1 FF, Label in: 100528, Label out: 3
    Time left: 125, Since: Thu Oct 21 15:29:26 2004
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 2 receiver 39082 protocol 0
    PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 17 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.1.13.1 (so-0/0/2.0) 17 pkts
    RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 17 pkts
    Explct route: 10.1.13.1
    Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
    LSPname: R1-to-R6 , LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: 3
    Resv style: 1 FF, Label in: 100544, Label out: 3
    Time left: 147, Since: Thu Oct 21 15:39:33 2004
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 1 receiver 47963 protocol 0
    PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 4 pkts
```



```

Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 4 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 4 pkts
Explot route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

```

Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPName: R6-to-R1
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.36.1 10.1.13.1
    18 Oct 21 15:34:18 Selected as active path
    17 Oct 21 15:34:17 Record Route: 10.1.36.1 10.1.13.1
    16 Oct 21 15:34:17 Up
    15 Oct 21 15:34:17 Originate Call
    14 Oct 21 15:34:17 CSPF: computation result accepted
    [...Output truncated...]
  Created: Tue Oct 19 22:28:30 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPName: R1-to-R6 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -

```



```

Time left: 126, Since: Thu Oct 21 15:44:25 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47963 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 4 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient

Record route: 10.1.13.1 10.1.36.1 <self>

Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Outputs 1 and 3 from ingress router R1 and egress router R6 show that the LSP is now traversing the network along the expected path, from R1 through R3 to R6, and the reverse LSP, from R6 through R3 to R1. In addition, Sample Output 2 from transit router R3 shows that there are two transit LSP sessions, one from R1 to R6, and the other from R6 to R1.

Checklist for Checking the RSVP Layer

Problem

Description: This checklist provides the steps and commands for checking the Resource Reservation Protocol (RSVP) layer of the layered Multiprotocol Label Switching (MPLS) model. The checklist provides links to an overview of the RSVP layer and more detailed information about the commands used to investigate the problem.

[Table 55 on page 2055](#) provides commands for checking the RSVP layer.

Solution

Table 55: Checklist for Checking the RSVP Layer

Tasks	Command or Action
“Checking the RSVP Layer” on page 2056	
1. Verify the LSP on page 2059	show mpls lsp extensive
2. Verify RSVP Sessions on page 2060	show rsvp session
3. Verify RSVP Neighbors on page 2063	show rsvp neighbor

Table 55: Checklist for Checking the RSVP Layer (*continued*)

Tasks	Command or Action
4. Verify RSVP Interfaces on page 2064	show rsvp interface
5. Verify the RSVP Protocol Configuration on page 2066	show configuration protocols rsvp
6. Take Appropriate Action on page 2067	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre>[edit] edit protocols rsvp [edit protocols rsvp] show set interface <i>type-fpc/pic/port</i> show commit</pre>
7. Verify the LSP Again on page 2069	show mpls lsp extensive

Checking the RSVP Layer

Purpose

After you have configured the label-switched path (LSP), issued the **show mpls lsp extensive** command, and determined that there is an error, you might find that the error is not in the physical, data link, or Internet Protocol (IP) and interior gateway protocol (IGP) layers. Continue investigating the problem at the RSVP layer of the network.

[Figure 158 on page 2057](#) illustrates the RSVP layer of the layered MPLS model.

Figure 158: Checking the RSVP Layer

BGP Layer	tracroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> tracroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
IGP and IP Layers Functioning	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive "JUNOS Interfaces Operations Guide"
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

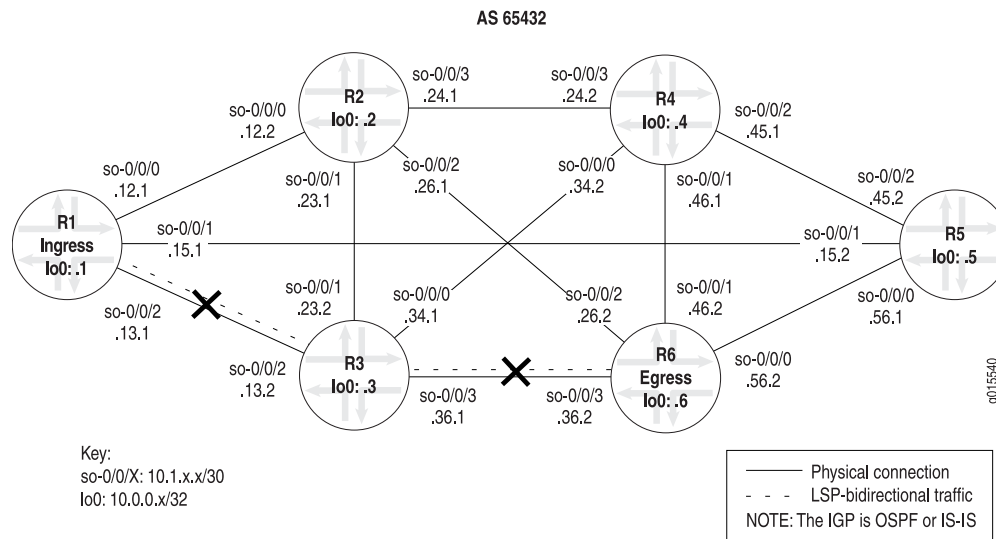
g015546

With this layer, you check that dynamic RSVP signaling is occurring as expected, neighbors are connected, and interfaces are configured correctly for RSVP. Check the ingress, egress, and transit routers.

If the network is not functioning at this layer, the LSP does not work as configured.

Figure 159 on page 2058 illustrates the MPLS network used in this topic.

Figure 159: MPLS Network Broken at the RSVP Layer



The network shown in [Figure 159 on page 2058](#) is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface. The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, the LSP is down without a path in either direction, from **R1** to **R6** or from **R6** to **R1**.

The crosses shown in [Figure 159 on page 2058](#) indicate where the LSP is broken. Some possible reasons the LSP is broken might include that dynamic RSVP signaling is not occurring as expected, neighbors are not connected, or interfaces are incorrectly configured for RSVP.

In the network in [Figure 159 on page 2058](#), a configuration error on transit router **R3** prevents the LSP from traversing the network as expected.

To check the RSVP layer, follow these steps:

1. [Verify the LSP | 2059](#)
2. [Verify RSVP Sessions | 2060](#)
3. [Verify RSVP Neighbors | 2063](#)
4. [Verify RSVP Interfaces | 2064](#)
5. [Verify the RSVP Protocol Configuration | 2066](#)
6. [Take Appropriate Action | 2067](#)
7. [Verify the LSP Again | 2069](#)

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To determine whether the LSP is up, enter the following command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0,  LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    2 Oct 27 15:06:05 10.1.13.2:  No Route toward dest [4 times]
    1 Oct 27 15:05:56 Originate Call
  Created: Wed Oct 27 15:05:55 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
```



```

Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0,  LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 22 second(s).
    1 Oct 27 14:59:12  CSPF failed: no route toward 10.0.0.1 [4 times]
  Created: Wed Oct 27 14:57:44 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output shows that the LSP is down in both directions, from **R1** to **R6**, and from **R6** to **R1**. The output from **R1** shows that **R1** is using a no-cspf LSP since it tried to originate the call without being able to reach the destination. The output from **R6** shows that the Constrained Shortest Path First (CSPF) algorithm failed, resulting in no route to destination **10.0.0.1**.

Verify RSVP Sessions

Purpose

When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, the LSP does not work as configured.

Action

To verify currently active RSVP sessions, enter the following command from the ingress, transit, and egress routers:


```
user@host> show rsvp session
```

Sample Output 1

```
user@R1> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```


Sample Output 2

```

user@R1> show rsvp session
Ingress RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
10.0.0.6     10.0.0.1     Up    1 1  FF      -    100768 R1-to-R6
Total 1 displayed, Up 1 , Down 0

Egress RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
10.0.0.1     10.0.0.6     Up    0 1  FF      3      -    R6-to-R1
Total 1 displayed, Up 1 , Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 2 sessions
To          From          State Rt  Style Labelin Labelout LSPname
10.0.0.1     10.0.0.6     Up    1 1  FF    100784      3    R6-to-R1
10.0.0.6     10.0.0.1     Up    1 1  FF    100768      3    R1-to-R6
Total 2 displayed, Up 2 , Down 0

user@R6> show rsvp session
Ingress RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
10.0.0.1     10.0.0.6     Up    1 1  FF      -    100784 R6-to-R1
Total 1 displayed, Up 1 , Down 0

Egress RSVP: 1 sessions
To          From          State Rt  Style Labelin Labelout LSPname
10.0.0.6     10.0.0.1     Up    0 1  FF      3      -    R1-to-R6
Total 1 displayed, Up 1 , Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```


Meaning

Sample Output 1 from all routers shows that no RSVP sessions were successfully created, even though the LSP **R6-to-R1** is configured. Continue investigating the problem in [“Verify RSVP Neighbors” on page 2063](#).

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the output from the ingress, transit, and egress routers when the RSVP configuration is correct, and the LSP is traversing the network as configured. **R1** and **R6** both show an ingress and egress RSVP session, with the LSP **R1-to-R6**, and the reverse LSP **R6-to-R1**. Transit router **R3** shows two transit RSVP sessions.

Verify RSVP Neighbors

Purpose

Display a list of RSVP neighbors that were learned dynamically when exchanging RSVP packets. Once a neighbor is learned, it is never removed from the list of RSVP neighbors unless the RSVP configuration is removed from the router.

Action

To verify RSVP neighbors, enter the following command from the ingress, transit, and egress routers:

```
user@host> show rsvp neighbor
```

Sample Output 1

```
user@R1> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.2         10  1/0      9:22          9    64/64    32

user@R3> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.1         0  1/0     28:20          9   190/190   41
10.1.36.2        16:50 1/1     15:37          9   105/78   38

user@R6> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.36.1        17:30 1/1     16:15          9   104/78   39
```


Sample Output 2

```

user@R3> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.1         5  1/0      9:14         9    63/63    33
10.1.36.2         5  1/0      9:05         9    62/62    32

user@R6> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.36.1         5  1/0      8:54         9    61/61    32

```

Meaning

Sample Output 1 shows that **R1** and **R6** have one RSVP neighbor each, **R3**. However, the values in the **Up/Dn** field are different. **R1** has a value of **1/0** and **R6** has a value of **1/1**, indicating that **R1** is an active neighbor with **R3**, but **R6** is not. When the up count is one more than the down count, the neighbor is active; if the values are equal, the neighbor is down. The values for **R6** are equal, **1/1**, indicating that the neighbor **R3** is down.

Transit router **R3** knows about two neighbors, **R1** and **R6**. The **Up/Dn** field indicates that **R1** is an active neighbor and **R6** is down. At this point it is not possible to determine if the problem resides with **R3** or **R6**, because both neighbors are not active. Continue investigating the problem in [“Verify RSVP Interfaces” on page 2064](#).

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the correct neighbor relationship between transit router **R3** and egress router **R6**. The **Up/Dn** field shows the up count to be one more than the down count, **1/0**, indicating that the neighbors are active.

Verify RSVP Interfaces

Purpose

Display the status of each interface on which RSVP is enabled to determine where the configuration error occurred.

Action

To verify the status of RSVP interfaces, enter the following command from the ingress, transit, and egress routers:

```

user@host> show rsvp interface

```


Sample Output 1

```
user@R1> show rsvp interface
```

```
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R3> show rsvp interface
```

```
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
<<< Missing interface so-0/0/3.0
```

```
user@R6> show rsvp interface
```

```
RSVP interface: 4 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

Sample Output 2

```
user@R1> show rsvp interface
```

```
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps


```
so-0/0/2.0 Up 1 100% 155.52Mbps 155.52Mbps 0bps 0bps
```

```
user@R3> show rsvp interface
```

```
Rsvp interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R6> show rsvp interface
```

```
Rsvp interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning

Sample Output 1 shows that even though each router has interfaces that are up and have RSVP active, there are no reservations (**Active resv**) on any of the routers. In this example, we would expect at least one reservation on the ingress and egress routers, and two reservations on the transit router.

In addition, interface **so-0/0/3** on transit router **R3** is not included in the configuration. The inclusion of this interface is critical to the success of the LSP.

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the relevant interfaces with active reservations.

Verify the RSVP Protocol Configuration

Purpose

After you have checked RSVP sessions, interfaces, neighbors, and determined that there might be a configuration error, verify the RSVP protocol configuration.

Action

To verify the RSVP configuration, enter the following command from the ingress, transit, and egress routers:


```
user@host> show configuration protocols rsvp
```

Sample Output

```
user@R1> show configuration protocols rsvp
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}

user@R3> show configuration protocols rsvp
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0; <<< Missing interface so-0/0/3.0
interface fxp0.0 {
    disable;
}

user@R6> show configuration protocols rsvp
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;
interface fxp0.0 {
    disable;
}
```

Meaning

The sample output shows that **R3** has interface **so-0/0/3.0** missing from the RSVP protocol configuration. This interface is critical for the correct functioning of the LSP.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, an interface is missing from the configuration of router R3.

Solution

To correct the error in this example, follow these steps:

1. Include the missing interface in the configuration of transit router R3:

```
user@R3> edit
user@R3# edit protocols rsvp
[edit protocols rsvp]
user@R3# show
user@R3# set interface so-0/0/3.0
```

2. Verify and commit the configuration:

```
[edit protocols rsvp]
user@R3# show
user@R3# commit
```

Sample Output

```
user@R3> edit
Entering configuration mode

[edit]
user@R3# edit protocols rsvp

[edit protocols rsvp]
user@R3# show
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0; <<< Missing interface so-0/0/3.0
interface fxp0.0 {
    disable;
}
[edit protocols rsvp]
user@R3# set interface so-0/0/3.0

[edit protocols rsvp]
user@R3# show
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
```



```

}
interface so-0/0/3.0; <<< Interface now included in the configuration

[edit protocols rsvp]
user@R3# commit
commit complete

```

Meaning

The sample output shows that the missing interface **so-0/0/3.0** on transit router **R3** is now correctly included at the **[edit protocols rsvp]** hierarchy level. This results in the possibility that the LSP might come up.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the MPLS layer has been resolved.

Action

To verify the LSP again, enter the following command on the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
      5 Oct 27 15:28:57 Selected as active path

```



```

    4 Oct 27 15:28:57   Record Route: 10.1.13.2 10.1.36.2
    3 Oct 27 15:28:57 Up
    2 Oct 27 15:28:44 10.1.13.2: No Route toward dest[35 times]
    1 Oct 27 15:05:56 Originate Call
    Created: Wed Oct 27 15:05:56 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 136, Since: Wed Oct 27 15:29:20 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39092 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 6 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 2

```

user@R3> show mpls lsp extensive

Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```


Transit LSP: 2 sessions

10.0.0.1

From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
 LSPname: R6-to-R1, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100672, Label out: 3
 Time left: 152, Since: Wed Oct 27 15:16:39 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 39092 protocol 0
 PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 7 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.13.1 (so-0/0/2.0) 7 pkts
 RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 7 pkts
 Explct route: 10.1.13.1
 Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
 LSPname: R1-to-R6, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 3
 Resv style: 1 FF, Label in: 100656, Label out: 3
 Time left: 129, Since: Wed Oct 27 14:53:14 2004
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 47977 protocol 0
 PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 40 pkts
 Adspec: received MTU 1500 sent MTU 1500
 PATH sentto: 10.1.36.2 (so-0/0/3.0) 7 pkts
 RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 7 pkts
 Record route: 10.1.13.1 <self> 10.1.36.2

Total 2 displayed, Up 2, Down 0

Sample Output 3

user@R6> show mpls lsp extensive

Ingress LSP: 1 sessions

10.0.0.1


```

From: 10.0.0.6, State: Up, ActiveRoute: 1 , LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.36.1 10.1.13.1
      6 Oct 27 15:22:06 Selected as active path
      5 Oct 27 15:22:06 Record Route: 10.1.36.1 10.1.13.1
      4 Oct 27 15:22:06 Up
      3 Oct 27 15:22:06 Originate Call
      2 Oct 27 15:22:06 CSPF: computation result accepted
      1 Oct 27 15:21:36 CSPF failed: no route toward 10.0.0.1[50 times]
  Created: Wed Oct 27 14:57:45 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 119, Since: Wed Oct 27 15:21:43 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 47977 protocol 0
  PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 7 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Verify the LSP

Purpose

Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action

To determine whether the LSP is up, enter the following command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0,  LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    2 Oct 27 15:06:05 10.1.13.2:  No Route toward dest [4 times]
    1 Oct 27 15:05:56 Originate Call
  Created: Wed Oct 27 15:05:55 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```



```

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0,  LSPname: R6-to-R1
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Will be enqueued for recomputation in 22 second(s).
    1 Oct 27 14:59:12  CSPF failed: no route toward 10.0.0.1 [4 times]
  Created: Wed Oct 27 14:57:44 2004
Total 1 displayed, Up 0,  Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output shows that the LSP is down in both directions, from **R1** to **R6**, and from **R6** to **R1**. The output from **R1** shows that **R1** is using a no-cspf LSP since it tried to originate the call without being able to reach the destination. The output from **R6** shows that the Constrained Shortest Path First (CSPF) algorithm failed, resulting in no route to destination **10.0.0.1**.

Verify RSVP Sessions

Purpose

When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, the LSP does not work as configured.

Action

To verify currently active RSVP sessions, enter the following command from the ingress, transit, and egress routers:

```
user@host> show rsvp session
```

Sample Output 1

```
user@R1> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```



```
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
user@R1> show rsvp session
Ingress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.0.0.6     10.0.0.1     Up    1 1 FF      -    100768  R1-to-R6
Total 1 displayed,  Up 1 , Down 0

Egress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.0.0.1     10.0.0.6     Up    0 1 FF      3      -  R6-to-R1
Total 1 displayed,  Up 1 , Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP:  2 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.0.0.1     10.0.0.6     Up    1 1 FF  100784      3  R6-to-R1
10.0.0.6     10.0.0.1     Up    1 1 FF  100768      3  R1-to-R6
Total 2 displayed,  Up 2 , Down 0

user@R6> show rsvp session
Ingress RSVP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.0.0.1     10.0.0.6     Up    1 1 FF      -    100784  R6-to-R1
Total 1 displayed,  Up 1 , Down 0
```



```

Egress RSVP: 1 sessions
To                From                State Rt  Style Labelin Labelout LSPname
10.0.0.6          10.0.0.1          Up    0 1 FF      3      -  R1-to-R6
Total 1 displayed,  Up 1 , Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 from all routers shows that no RSVP sessions were successfully created, even though the LSP **R6-to-R1** is configured. Continue investigating the problem in [“Verify RSVP Neighbors” on page 2063](#).

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the output from the ingress, transit, and egress routers when the RSVP configuration is correct, and the LSP is traversing the network as configured. **R1** and **R6** both show an ingress and egress RSVP session, with the LSP **R1-to-R6**, and the reverse LSP **R6-to-R1**. Transit router **R3** shows two transit RSVP sessions.

Verify RSVP Neighbors

Purpose

Display a list of RSVP neighbors that were learned dynamically when exchanging RSVP packets. Once a neighbor is learned, it is never removed from the list of RSVP neighbors unless the RSVP configuration is removed from the router.

Action

To verify RSVP neighbors, enter the following command from the ingress, transit, and egress routers:

```
user@host> show rsvp neighbor
```

Sample Output 1

```

user@R1> show rsvp neighbor
RSVP neighbor: 1 learned
Address                Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.2              10  1/0      9:22         9     64/64    32

```



```

user@R3> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.1         0  1/0    28:20         9    190/190   41
10.1.36.2        16:50 1/1    15:37         9    105/78   38

user@R6> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.36.1        17:30 1/1    16:15         9    104/78   39

```

Sample Output 2

```

user@R3> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.13.1         5  1/0     9:14         9     63/63   33
10.1.36.2         5  1/0     9:05         9     62/62   32

user@R6> show rsvp neighbor
RSVP neighbor: 1 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx MsgRcvd
10.1.36.1         5  1/0     8:54         9     61/61   32

```

Meaning

Sample Output 1 shows that **R1** and **R6** have one RSVP neighbor each, **R3**. However, the values in the **Up/Dn** field are different. **R1** has a value of **1/0** and **R6** has a value of **1/1**, indicating that **R1** is an active neighbor with **R3**, but **R6** is not. When the up count is one more than the down count, the neighbor is active; if the values are equal, the neighbor is down. The values for **R6** are equal, **1/1**, indicating that the neighbor **R3** is down.

Transit router **R3** knows about two neighbors, **R1** and **R6**. The **Up/Dn** field indicates that **R1** is an active neighbor and **R6** is down. At this point it is not possible to determine if the problem resides with **R3** or **R6**, because both neighbors are not active. Continue investigating the problem in [“Verify RSVP Interfaces” on page 2064](#).

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the correct neighbor relationship between transit router **R3** and egress router **R6**. The **Up/Dn** field shows the up count to be one more than the down count, **1/0**, indicating that the neighbors are active.

Verify RSVP Interfaces

Purpose

Display the status of each interface on which RSVP is enabled to determine where the configuration error occurred.

Action

To verify the status of RSVP interfaces, enter the following command from the ingress, transit, and egress routers:

```
user@host> show rsvp interface
```

Sample Output 1

```
user@R1> show rsvp interface
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R3> show rsvp interface
RSVP interface: 3 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
<<< Missing interface so-0/0/3.0
```

```
user@R6> show rsvp interface
RSVP interface: 4 active
```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps


```
so-0/0/2.0 Up          0  100% 155.52Mbps 155.52Mbps 0bps      0bps
so-0/0/3.0 Up          0  100% 155.52Mbps 155.52Mbps 0bps      0bps
```

Sample Output 2

```
user@R1> show rsvp interface
```

```
Rsvp interface: 3 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R3> show rsvp interface
```

```
Rsvp interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R6> show rsvp interface
```

```
Rsvp interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning

Sample Output 1 shows that even though each router has interfaces that are up and have RSVP active, there are no reservations (**Active resv**) on any of the routers. In this example, we would expect at least one reservation on the ingress and egress routers, and two reservations on the transit router.

In addition, interface **so-0/0/3** on transit router **R3** is not included in the configuration. The inclusion of this interface is critical to the success of the LSP.

In contrast to Sample Output 1 and to illustrate correct output, Sample Output 2 shows the relevant interfaces with active reservations.

Verify the RSVP Protocol Configuration

Purpose

After you have checked RSVP sessions, interfaces, neighbors, and determined that there might be a configuration error, verify the RSVP protocol configuration.

Action

To verify the RSVP configuration, enter the following command from the ingress, transit, and egress routers:

```
user@host> show configuration protocols rsvp
```

Sample Output

```
user@R1> show configuration protocols rsvp
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}

user@R3> show configuration protocols rsvp
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0; <<< Missing interface so-0/0/3.0
interface fxp0.0 {
    disable;
}

user@R6> show configuration protocols rsvp
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
```



```
interface so-0/0/3.0;
interface fxp0.0 {
    disable;
}
```

Meaning

The sample output shows that **R3** has interface **so-0/0/3.0** missing from the RSVP protocol configuration. This interface is critical for the correct functioning of the LSP.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, an interface is missing from the configuration of router R3.

Solution

To correct the error in this example, follow these steps:

1. Include the missing interface in the configuration of transit router R3:

```
user@R3> edit
user@R3# edit protocols rsvp
[edit protocols rsvp]
user@R3# show
user@R3# set interface so-0/0/3.0
```

2. Verify and commit the configuration:

```
[edit protocols rsvp]
user@R3# show
user@R3# commit
```

Sample Output

```
user@R3> edit
Entering configuration mode

[edit]
user@R3# edit protocols rsvp
```



```

[edit protocols rsvp]
user@R3# show
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0; <<< Missing interface so-0/0/3.0
interface fxp0.0 {
    disable;
}
[edit protocols rsvp]
user@R3# set interface so-0/0/3.0

[edit protocols rsvp]
user@R3# show
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}
interface so-0/0/3.0; <<< Interface now included in the configuration

[edit protocols rsvp]
user@R3# commit
commit complete

```

Meaning

The sample output shows that the missing interface **so-0/0/3.0** on transit router **R3** is now correctly included at the **[edit protocols rsvp]** hierarchy level. This results in the possibility that the LSP might come up.

Verify the LSP Again

Purpose

After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the MPLS layer has been resolved.

Action

To verify the LSP again, enter the following command on the ingress, transit, and egress routers:


```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up,  ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.13.2 10.1.36.2
    5 Oct 27 15:28:57 Selected as active path
    4 Oct 27 15:28:57  Record Route: 10.1.13.2 10.1.36.2
    3 Oct 27 15:28:57 Up
    2 Oct 27 15:28:44 10.1.13.2: No Route toward dest[35 times]
    1 Oct 27 15:05:56 Originate Call
    Created: Wed Oct 27 15:05:56 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 136, Since: Wed Oct 27 15:29:20 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39092 protocol 0
  PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 6 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0
```



```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Sample Output 2

```
user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1
  From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
  LSPname: R6-to-R1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100672, Label out: 3
  Time left: 152, Since: Wed Oct 27 15:16:39 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 39092 protocol 0
  PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 7 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.1.13.1 (so-0/0/2.0) 7 pkts
  RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 7 pkts
  Explct route: 10.1.13.1
  Record route: 10.1.36.2 <self> 10.1.13.1

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100656, Label out: 3
```



```

Time left: 129, Since: Wed Oct 27 14:53:14 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47977 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 40 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 7 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 7 pkts
Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

```

Sample Output 3

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Up, ActiveRoute: 1 , LSPname: R6-to-R1
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.1.36.1 10.1.13.1
  6 Oct 27 15:22:06 Selected as active path
  5 Oct 27 15:22:06 Record Route: 10.1.36.1 10.1.13.1
  4 Oct 27 15:22:06 Up
  3 Oct 27 15:22:06 Originate Call
  2 Oct 27 15:22:06 CSPF: computation result accepted
  1 Oct 27 15:21:36 CSPF failed: no route toward 10.0.0.1[50 times]
  Created: Wed Oct 27 14:57:45 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: R1-to-R6, LSPpath: Primary

```



```

Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 119, Since: Wed Oct 27 15:21:43 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 47977 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 7 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Checklist for Determining LSP Status

Purpose

This checklist provides the steps and commands to verify the state of a label-switched path (LSP) in an MPLS network. The checklist includes links to more detailed information about the commands to verify the LSP and supporting protocols. [Table 56 on page 2087](#) provides commands for determining the LSP state.

Action

Table 56: Checklist for Determining the LSP State

Tasks	Command or Action
“Determining LSP Status” on page 2285	
1. Check the Status of the LSP on page 2286	show mpls lsp

Table 56: Checklist for Determining the LSP State *(continued)*

Tasks	Command or Action
2. Display Extensive Status About the LSP on page 2287	<code>show mpls lsp extensive</code>
“Determining LSP Statistics” on page 2088	<code>show rsvp session detail</code>

Determining LSP Statistics

Purpose

Display detailed information about RSVP objects to assist the diagnosis of an LSP problem.

Action

To verify RSVP objects, enter the following Junos OS CLI operational mode command:

```
user@host> show rsvp session detail
```

Sample Output

```
user@R1> show rsvp session detail
Ingress RSVP: 1 sessions

10.0.0.6
  From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
  LSPname: R1-to-R6 , LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100064
  Resv style: 1 FF, Label in: -, Label out: 100064
  Time left: -, Since: Tue Aug 17 12:22:52 2004
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 12 receiver 44251 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 10.1.13.2 (so-0/0/2.0) 182 pkts
  RESV rcvfrom: 10.1.13.2 (so-0/0/2.0) 159 pkts
  Explct route: 10.1.13.2 10.1.36.2
  Record route: <self> 10.1.13.2 10.1.36.2
```



```
Total 1 displayed, Up 1, Down 0
```

```
Egress RSVP: 1 sessions
```

10.0.0.1

```
From: 10.0.0.6 , LSPstate: Up, ActiveRoute: 0
LSPname: R6-to-R1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 135, Since: Tue Aug 17 12:23:14 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39024 protocol 0
PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 158 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.56.2 10.1.15.2 <self>
Total 1 displayed, Up 1, Down 0
```

```
Transit RSVP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

Meaning

The sample output shows that there is one ingress and one egress RSVP session. The ingress session has a source address of **10.0.0.1 (R1)**, and the session is up, with one active route. The LSP name is **R1-to-R6** and it is the primary path for the LSP.

The recovery label (**100064**) is sent by a graceful restart router to its neighbor to recover a forwarding state. It is probably the old label that the router advertised before it went down.

This session is using the fixed filter (**FF**) reservation style (**Resv style**). Since this is an ingress router, there is no inbound label. The outbound label (provided by the next downstream router) is **100064**.

The **Time Left** field provides the number of seconds remaining in the RSVP session, and the **Tspec** object provides information about the controlled load rate (**rate**) and maximum burst size (**peak**), an infinite value (**Infbps**) for the guaranteed delivery option, and the indication that packets smaller than 20 bytes are treated as 20 bytes, while packets larger than 1500 bytes are treated as 1500 bytes.

The port number is the IPv4 tunnel ID, while the sender/receiver port number is the LSP ID. The IPv4 tunnel ID is unique for the life of the LSP, while the sender/receiver LSP ID can change, for example, with an SE style reservation.

The **PATH rcvfrom** field includes the source of the path message. Since this is the ingress router, the local client originated the path message.

The **PATH sentto** field includes the path message destination (**10.1.13.2**) and outgoing interface (**so-0/0/2.0**). The **RESV rcvfrom** field includes both the source of the Resv message received (**10.1.13.2**) and the incoming interface (**so-0/0/2.0**).

The RSVP explicit route and the route record values are identical: **10.1.13.2** and **10.1.36.2**. In most cases, the explicit route and the record route values are identical. Differences indicate that some path rerouting has occurred, typically during Fast-Reroute.

The **Total** fields indicate the total number of ingress, egress, and transit RSVP sessions, with the total being equal to the sum of the up and down sessions. In this example, there is one ingress session, one egress session, and no transit RSVP sessions.

Checklist for Verifying LSP Use

Purpose

This checklist provides the steps and commands to verify the use of the LSP in an MPLS network. The checklist includes links to more detailed information about verifying the LSP on the ingress and transit routers in the network.

This checklist describes how to verify the availability and valid use of a label-switched path (LSP) in your network.

Action

Table 57: Checklist for Verifying LSP Use

Tasks	Command or Action
"Verifying LSP Use in Your Network" on page 2090	
"Verifying an LSP on the Ingress Router" on page 2091	<code>show route table inet.3</code>
"Verifying an LSP on a Transit Router" on page 2093	<code>show route table mpls.0</code>

[Table 57 on page 2090](#) provides commands for verifying LSP use.

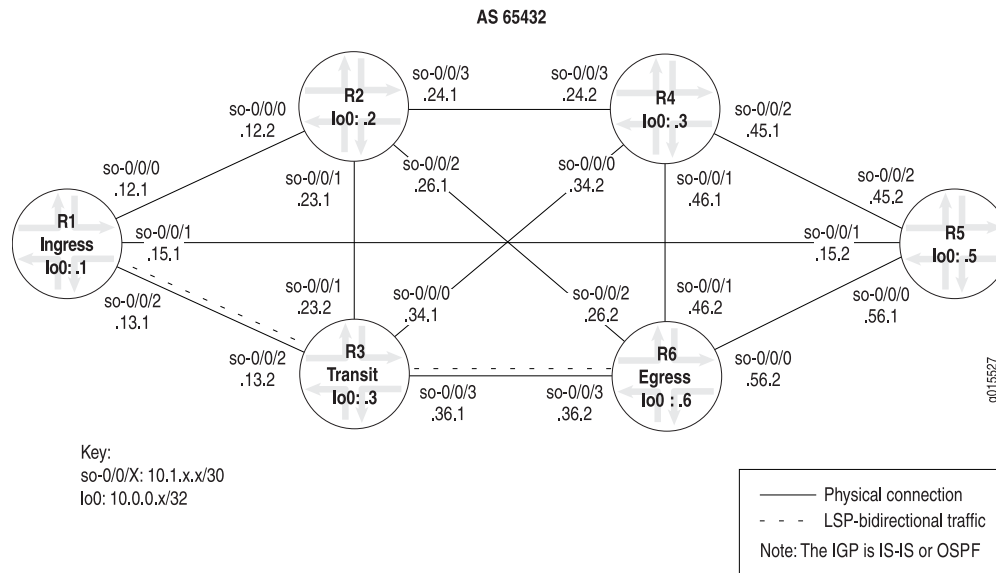
Verifying LSP Use in Your Network

Purpose

When you verify the valid use of an LSP on the ingress and transit routers in your network, you can determine if there is a problem with Multiprotocol Label Switching (MPLS) in your network.

[Figure 160 on page 2091](#) describes the example network used in this topic.

Figure 160: MPLS Topology for Verifying LSP Use



The MPLS network in [Figure 160 on page 2091](#) illustrates a router-only network with SONET interfaces that consist of the following components:

- A full-mesh interior Border Gateway Protocol (IBGP) topology, using AS 65432
- MPLS and Resource Reservation Protocol (RSVP) enabled on all routers
- A **send-statics** policy on routers R1 and R6 that allows a new route to be advertised into the network
- An LSP between routers R1 and R6

The network shown in [Figure 160 on page 2091](#) is a Border Gateway Protocol (BGP) full-mesh network. Since route reflectors and confederations are not used to propagate BGP learned routes, each router must have a BGP session with every other router running BGP.

To verify LSP use in your network, follow these steps:

1. [Verifying an LSP on the Ingress Router | 2091](#)
2. [Verifying an LSP on a Transit Router | 2093](#)

Verifying an LSP on the Ingress Router

Purpose

You can verify the availability of an LSP when it is up by examining the **inet.3** routing table on the ingress router. The **inet.3** routing table contains the host address of each LSP's egress router. This routing table

is used on ingress routers to route BGP packets to the destination egress router. BGP uses the **inet.3** routing table on the ingress router to help resolve next-hop addresses.

Action

To verify an LSP on an ingress router, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show route table inet.3
```

Sample Output

```
user@R1> show route table inet.3
```

```
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32          *[RSVP/7] 4w0d 22:40:57, metric 20
                    > via so-0/0/2.0, label-switched-path R1-to-R6
```

Meaning

The sample output shows the **inet.3** routing table. By default, only BGP and MPLS virtual private networks (VPNs) can use the **inet.3** route table to resolve next-hop information. One destination is listed in the route table, **10.0.0.6**. This destination (**10.0.0.6**) is signaled by RSVP, and is the current active path, as indicated by the asterisk (*). The protocol preference for this route is **7**, and the metric associated with it is **20**. The label-switched path is **R1-to-R6**, through interface **so-0/0/2.0**, which is the physical next-hop transit interface.

Typically, the penultimate router in the LSP either pops the packet's label or changes the label to a value of 0. If the penultimate router pops the top label and an IPv4 packet is underneath, the egress router routes the IPv4 packet, consulting the IP routing table **inet.0** to determine how to forward the packet. If another type of label (such as one created by Label Distribution Protocol (LDP) tunneling or VPNs, but not IPv4) is underneath the top label, the egress router does not examine the **inet.0** routing table. Instead, it examines the **mpls.0** routing table for forwarding decisions.

If the penultimate router changes the packet's label to a value of 0, the egress router strips off the 0 label, indicating that an IPv4 packet follows. The packet is examined by the **inet.0** routing table for forwarding decisions.

When a transit or egress router receives an MPLS packet, information in the MPLS forwarding table is used to determine the next transit router in the LSP or whether this router is the egress router.

When BGP resolves a next-hop prefix, it examines both the **inet.0** and **inet.3** routing tables, seeking the next hop with the lowest preference; for example, RSVP preference 7 is preferred over OSPF preference 10. The RSVP signaled LSP is used to reach the BGP next hop. This is the default when the BGP next hop equals the LSP egress address. Once the BGP next hop is resolved through an LSP, the BGP traffic uses the LSP to forward BGP transit traffic.

Verifying an LSP on a Transit Router

Purpose

You can verify the availability of an LSP when it is up by examining the **mpls.0** routing table on a transit router. MPLS maintains the **mpls.0** routing table, which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

Action

To verify an LSP on a transit router, enter the following Junos OS CLI operational mode command:

```
user@host> show route table mpls.0
```

Sample Output

```
user@R3> show route table mpls.0
```

```
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          * [MPLS/0] 7w3d 22:20:56, metric 1
            Receive
1          * [MPLS/0] 7w3d 22:20:56, metric 1
            Receive
2          * [MPLS/0] 7w3d 22:20:56, metric 1
            Receive
100064     * [RSVP/7] 2w1d 04:17:36, metric 1
            > via so-0/0/3.0, label-switched-path R1-to-R6
100064 (S=0) * [RSVP/7] 2w1d 04:17:36, metric 1
            > via so-0/0/3.0, label-switched-path R1-to-R6
```

Meaning

The sample output from transit router **R3** shows route entries in the form of MPLS label entries, indicating that there is only one active route, even though there are five active entries.

The first three MPLS labels are reserved MPLS labels defined in RFC 3032. Packets received with these label values are sent to the Routing Engine for processing. Label 0 is the IPv4 explicit null label. Label 1 is the MPLS equivalent of the IP Router Alert label and Label 2 is the IPv6 explicit null label.

The two entries with the **100064** label are for the same LSP, **R1-to-R6**. There are two entries because the stack values in the MPLS header may be different. The second entry, **100064 (S=0)**, indicates that the stack depth is not 1 and additional label values are included in the packet. In contrast, the first entry of **100064** has an inferred S=1 which indicates a stack depth of 1 and makes it the last label in the packet. The dual entry indicates that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

The incoming label is the MPLS header of the MPLS packet, and is assigned by RSVP to the upstream neighbor. Juniper Networks routers dynamically assign labels for RSVP traffic-engineered LSPs in the range from 100,000 through 1,048,575.

The router assigns labels starting at label 100,000, in increments of 16. The sequence of label assignments is 100,000, 100,016, 100,032, 100,048, and so on. At the end of the assigned labels, the label numbers start over at 100001, incrementing in units of 16. Juniper Networks reserves labels for various purposes. [Table 58 on page 2094](#) lists the various label range allocations for incoming labels.

Table 58: MPLS Label Range Allocations

Incoming Label	Status
0 through 15	Reserved by IETF
16 through 1023	Reserved for static LSP assignment
1024 through 9999	Reserved for internal use (for example, CCC labels)
10,000 through 99,999	Reserved for static LSP assignment
100,000 through 1,048,575	Reserved for dynamic label assignment

Verifying an LSP on the Ingress Router

Purpose

You can verify the availability of an LSP when it is up by examining the **inet.3** routing table on the ingress router. The **inet.3** routing table contains the host address of each LSP's egress router. This routing table is used on ingress routers to route BGP packets to the destination egress router. BGP uses the **inet.3** routing table on the ingress router to help resolve next-hop addresses.

Action

To verify an LSP on an ingress router, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show route table inet.3
```

Sample Output

```
user@R1> show route table inet.3
```

```
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32          *[RSVP/7] 4w0d 22:40:57, metric 20
                    > via so-0/0/2.0, label-switched-path R1-to-R6
```

Meaning

The sample output shows the **inet.3** routing table. By default, only BGP and MPLS virtual private networks (VPNs) can use the **inet.3** route table to resolve next-hop information. One destination is listed in the route table, **10.0.0.6**. This destination (**10.0.0.6**) is signaled by RSVP, and is the current active path, as indicated by the asterisk (*). The protocol preference for this route is **7**, and the metric associated with it is **20**. The label-switched path is **R1-to-R6**, through interface **so-0/0/2.0**, which is the physical next-hop transit interface.

Typically, the penultimate router in the LSP either pops the packet's label or changes the label to a value of 0. If the penultimate router pops the top label and an IPv4 packet is underneath, the egress router routes the IPv4 packet, consulting the IP routing table **inet.0** to determine how to forward the packet. If another type of label (such as one created by Label Distribution Protocol (LDP) tunneling or VPNs, but not IPv4) is underneath the top label, the egress router does not examine the **inet.0** routing table. Instead, it examines the **mpls.0** routing table for forwarding decisions.

If the penultimate router changes the packet's label to a value of 0, the egress router strips off the 0 label, indicating that an IPv4 packet follows. The packet is examined by the **inet.0** routing table for forwarding decisions.

When a transit or egress router receives an MPLS packet, information in the MPLS forwarding table is used to determine the next transit router in the LSP or whether this router is the egress router.

When BGP resolves a next-hop prefix, it examines both the **inet.0** and **inet.3** routing tables, seeking the next hop with the lowest preference; for example, RSVP preference 7 is preferred over OSPF preference 10. The RSVP signaled LSP is used to reach the BGP next hop. This is the default when the BGP next hop

equals the LSP egress address. Once the BGP next hop is resolved through an LSP, the BGP traffic uses the LSP to forward BGP transit traffic.

Verifying an LSP on a Transit Router

Purpose

You can verify the availability of an LSP when it is up by examining the **mpls.0** routing table on a transit router. MPLS maintains the **mpls.0** routing table, which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

Action

To verify an LSP on a transit router, enter the following Junos OS CLI operational mode command:

```
user@host> show route table mpls.0
```

Sample Output

```
user@R3> show route table mpls.0
```

```
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                * [MPLS/0] 7w3d 22:20:56, metric 1
                  Receive
1                * [MPLS/0] 7w3d 22:20:56, metric 1
                  Receive
2                * [MPLS/0] 7w3d 22:20:56, metric 1
                  Receive
100064            * [RSVP/7] 2w1d 04:17:36, metric 1
                  > via so-0/0/3.0, label-switched-path R1-to-R6
100064 (S=0)      * [RSVP/7] 2w1d 04:17:36, metric 1
                  > via so-0/0/3.0, label-switched-path R1-to-R6
```

Meaning

The sample output from transit router **R3** shows route entries in the form of MPLS label entries, indicating that there is only one active route, even though there are five active entries.

The first three MPLS labels are reserved MPLS labels defined in RFC 3032. Packets received with these label values are sent to the Routing Engine for processing. Label 0 is the IPv4 explicit null label. Label 1 is the MPLS equivalent of the IP Router Alert label and Label 2 is the IPv6 explicit null label.

The two entries with the **100064** label are for the same LSP, **R1-to-R6**. There are two entries because the stack values in the MPLS header may be different. The second entry, **100064 (S=0)**, indicates that the stack depth is not 1 and additional label values are included in the packet. In contrast, the first entry of **100064** has an inferred S=1 which indicates a stack depth of 1 and makes it the last label in the packet. The dual entry indicates that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

The incoming label is the MPLS header of the MPLS packet, and is assigned by RSVP to the upstream neighbor. Juniper Networks routers dynamically assign labels for RSVP traffic-engineered LSPs in the range from 100,000 through 1,048,575.

The router assigns labels starting at label 100,000, in increments of 16. The sequence of label assignments is 100,000, 100,016, 100,032, 100,048, and so on. At the end of the assigned labels, the label numbers start over at 100001, incrementing in units of 16. Juniper Networks reserves labels for various purposes. [Table 58 on page 2094](#) lists the various label range allocations for incoming labels.

Table 59: MPLS Label Range Allocations

Incoming Label	Status
0 through 15	Reserved by IETF
16 through 1023	Reserved for static LSP assignment
1024 through 9999	Reserved for internal use (for example, CCC labels)
10,000 through 99,999	Reserved for static LSP assignment
100,000 through 1,048,575	Reserved for dynamic label assignment

Verify That Load Balancing Is Working

Purpose

After configuring load balancing, check that traffic is load-balanced equally across paths. In this section, the command output reflects the load-balancing configuration of the example network shown in [“Load-Balancing Network Topology” on page 2101](#). The **clear** commands are used to reset LSP and interface counters to zero so that the values reflect the operation of the load-balancing configuration.

Action

To verify load balancing across interfaces and LSPs, use the following command on the ingress router:

```
user@host# show configuration
```

To verify load balancing across interfaces and LSPs, use the following commands on a transit router:

```
user@host# show route
user@host# show route forwarding-table
user@host# show mpls lsp statistics
user@host# monitor interface traffic
user@host# clear mpls lsp statistics
user@host# clear interface statistics
```

Sample Output

The following sample output is for the configuration on ingress router **R1**:

```
user@R1> show configuration | no-more
[...Output truncated...]
routing-options {
  [...Output truncated...]
  forwarding-table {
    export lbpp;
  }
}
[...Output truncated...]
policy-options {
  policy-statement lbpp {
    then {
      load-balance per-packet;
    }
  }
}
```

Meaning

The sample output for the **show configuration** command on ingress router **R1** shows that load balancing is correctly configured with the **lbpp** policy statement. Also, the **lbpp** policy is exported into the forwarding table at the **[edit routing-options]** hierarchy level.

Sample Output

The following sample output is from transit router R2:

```
user@R2> show route 192.168.0.1 terse

inet.0: 25 destinations, 27 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1   Metric 2  Next hop      AS path
* 192.168.0.1/32   O  10           3           so-0/0/1.0
                                >so-0/0/2.0

[...Output truncated...]
```

Meaning

The sample output for the **show route** command issued on transit router **R2** shows the two equal-cost paths (**so-0/0/1** and **so-0/0/2**) through the network to the loopback address to **R0 (192.168.0.1)**. Even though the right angle bracket (>) usually indicates the active route, in this instance it does not, as shown in the following four sample outputs.

Sample Output

The following sample output is from transit router R2:

```
user@R2> monitor interface traffic

R2                               Seconds: 65                               Time: 11:41:14

Interface  Link  Input packets  (pps)  Output packets  (pps)
so-0/0/0   Up    0              (0)    0              (0)
so-0/0/1   Up    126            (0)    164659         (2128)
so-0/0/2   Up    85219          (1004) 164598         (2128)
so-0/0/3   Up    0              (0)    0              (0)
fe-0/1/0   Up    328954         (4265) 85475          (1094)
fe-0/1/1   Up    0              (0)    0              (0)
fe-0/1/2   Up    0              (0)    0              (0)
fe-0/1/3   Up    0              (0)    0              (0)

[...Output truncated...]
```

Meaning

The sample output for the **monitor interface traffic** command issued on transit router **R2** shows that output traffic is evenly distributed across the two interfaces **so-0/0/1** and **so-0/0/2**.

Sample Output

The following sample output is from transit router R2:

```
user@R2> show mpls lsp statistics
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 5 sessions
To          From          State   Packets   Bytes LSPname
192.168.0.1 192.168.1.1 Up      87997    17951388 lsp1
192.168.0.1 192.168.1.1 Up      87997    17951388 lsp2
192.168.0.1 192.168.1.1 Up      87997    17951388 lsp3
192.168.0.1 192.168.1.1 Up      87997    17951388 lsp4
192.168.6.1 192.168.0.1 Up        0          0 r0-r1
Total 5 displayed, Up 5, Down 0
```

Meaning

The sample output for the **show mpls lsp statistics** command issued on transit router **R2** shows that output traffic is evenly distributed across the four LSPs configured on ingress router **R6**.

Sample Output

The following sample output is from transit router R2:

```
user@R2> show route forwarding-table destination 10.0.90.14
Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
10.0.90.12/30    user    0                ulst 262144 6
                  ucst 345  5 so-0/0/1.0
                  ucst 339  2 so-0/0/2.0
```

Meaning

The sample output for the **show route forwarding-table destination** command issued on transit router **R2** shows **ulst** in the **Type** field, which indicates that load balancing is working. The two unicast (**ucst**) entries in the **Type** field are the two next hops for the LSPs.

Sample Output

The following sample output is from transit router R2:

```
user@R2> show route forwarding-table | find mpls
```

Routing table: mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	38	1	
0	user	0		recv	37	3	
1	user	0		recv	37	3	
2	user	0		recv	37	3	
100112	user	0		Swap	100032		so-0/0/1.0
100128	user	0		Swap	100048		so-0/0/1.0
100144	user	0	10.0.12.13	Swap	100096		fe-0/1/0.0
100160	user	0		Swap	100112		so-0/0/2.0
100176	user	0		Swap	100128		so-0/0/2.0

Meaning

The sample output for the **show route forwarding-table | find mpls** command issued on transit router R2 shows the MPLS routing table that contains the labels received and used by this router to forward packets to the next-hop router. This routing table is used mostly on transit routers to route packets to the next router along an LSP. The first three labels in the **Destination** column (Label 0, Label 1, and Label 2) are automatically entered by MPLS when the protocol is enabled. These labels are reserved MPLS labels defined in RFC 3032. Label 0 is the IPv4 explicit null label. Label 1 is the MPLS equivalent of the IP Router Alert label, and Label 2 is the IPv6 explicit null label.

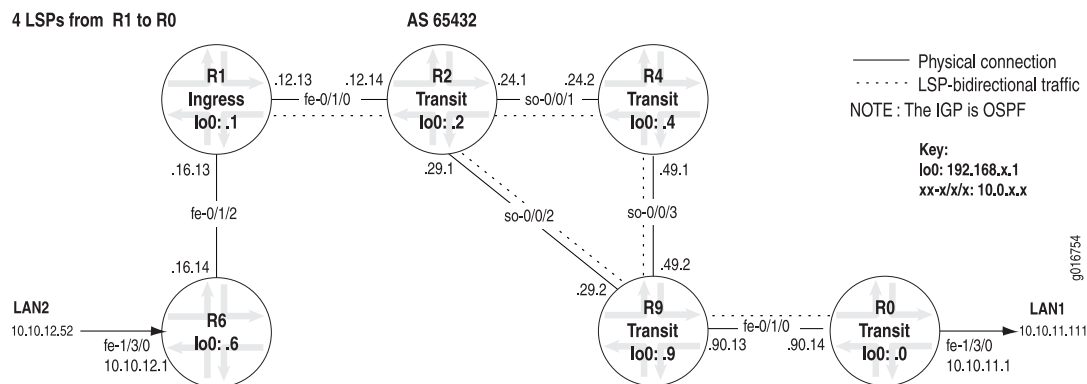
The remaining five labels in the **Destination** column are nonreserved labels that the router uses to forward traffic, and the last column **Netif**, shows the interfaces used to send the labeled traffic. For nonreserved labels, the second **Type** column shows the operation performed on matching packets. In this example, all non-reserved packets are swapped for outgoing packet labels. For example, packets with the label **100112** have their label swapped for **100032** before they are pushed out of interface **so-0/0/1.0**.

Example: Load-Balanced MPLS Network

When you configure several RSVP LSPs to the same egress router, the LSP with the lowest metric is selected and carries all traffic. If all of the LSPs have the same metric, one of the LSPs is selected at random and all traffic is forwarded over it. To distribute traffic equally across all LSPs, you can configure load balancing on the ingress or transit routers, depending on the type of load balancing configured.

Figure 161 on page 2102 illustrates an MPLS network with four LSPs configured to the same egress router (R0). Load balancing is configured on ingress router R1. The example network uses Open Shortest Path First (OSPF) as the interior gateway protocol (IGP) with OSPF area 0.0.0.0. An IGP is required for the Constrained Shortest Path First (CSPF) LSP, which is the default for the Junos OS. In addition, the example network uses a policy to create BGP traffic.

Figure 161: Load-Balancing Network Topology



The network shown in Figure 161 on page 2102 consists of the following components:

- A full-mesh interior BGP (IBGP) topology, using AS 65432
- MPLS and RSVP enabled on all routers
- A send-statics policy on routers R1 and R0 that allows a new route to be advertised into the network
- Four unidirectional LSPs between R1 and R0, and one reverse direction LSP between R0 and R1, which allows for bidirectional traffic
- Load balancing configured on ingress router R1

The network shown in Figure 161 on page 2102 is a BGP full-mesh network. Since route reflectors and confederations are not used to propagate BGP learned routes, each router must have a BGP session with every other router running BGP.

Router Configurations for the Load-Balanced MPLS Network

Purpose

The configurations in this topic are for the six routers in the example network illustrated in “Load-Balancing Network Topology” on page 2101.

Action

To display the configuration of a router, use the following Junos OS CLI operational mode command:


```
user@host> show configuration | no-more
```

Sample Output 1

The following configuration output is for edge router **R6**.

```
user@R6> show configuration | no-more
[...Output truncated...]
interfaces {
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.16.14/30;
            }
            family mpls; #MPLS enabled on relevant interfaces
        }
    }
    fe-1/3/0 {
        unit 0 {
            family inet {
                address 10.10.12.1/24;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.148/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.6.1/32;
            }
        }
    }
}
routing-options {
    static {
[...Output truncated...]
        router-id 192.168.6.1; #Manually configured RID
    }
}
```



```

    autonomous-system 65432; #Full mesh IBGP
  }
}
protocols {
  rsvp {
    interface fe-0/1/2.0;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    interface fe-0/1/2.0;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group internal {
      type internal;
      local-address 192.168.6.1;
      neighbor 192.168.1.1;
      neighbor 192.168.2.1;
      neighbor 192.168.4.1;
      neighbor 192.168.9.1;
      neighbor 192.168.0.1;
    }
  }
  ospf { #IGP enabled
    traffic-engineering;
    area 0.0.0.0 {
      interface fe-0/1/2.0;
      interface fe-1/3/0.0;
      interface lo0.0 {
        passive; #Ensures protocols do not run over this interface
      }
    }
  }
}

```

Sample Output 2

The following configuration output is for ingress router **R1**.


```

user@R1> show configuration | no-more
[...Output truncated...]
interfaces {
    fe-0/1/0 {
        unit 0 {
            family inet {
                address 10.0.12.13/30;
            }
            family mpls; #MPLS enabled on relevant interfaces
        }
    }
    fe-0/1/2 {
        unit 0 {
            family inet {
                address 10.0.16.13/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.143/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.1.1/32;
            }
        }
    }
}
routing-options {
    static {
        [...Output truncated...]
        route 100.100.1.0/24 reject; #Static route for send-statics policy
    }

    router-id 192.168.1.1; #Manually configured RID
    autonomous-system 65432; #Full mesh IBGP
    forwarding-table {
        export lbpp; #Routes exported to forwarding table
    }
}

```



```

}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface fe-0/1/2.0;
        interface fxp0.0 {
            disable;
        }
    }
}
mpls {
    label-switched-path lsp 1 { #First LSP
        to 192.168.0.1; # Destination of the LSP
        install 10.0.90.14/32 active; # The prefix is installed in the
        primary via-r4; # inet.0 routing table
    }
    label-switched-path lsp2 {
        to 192.168.0.1;
        install 10.0.90.14/32 active;
        primary via-r2;
    }
    label-switched-path lsp3 {
        to 192.168.0.1;
        install 10.0.90.14/32 active;
        primary via-r2;
    }
    label-switched-path lsp4 {
        to 192.168.0.1;
        install 10.0.90.14/32 active;
        primary via-r4;
    }
    path via-r2 { #Primary path to spread traffic across interfaces
        10.0.29.2 loose;
    }
    path via-r4 {
        10.0.24.2 loose;
    }
    interface fe-0/1/0.0;
    interface fe-0/1/2.0;
    interface fxp0.0 {
        disable;
    }
}
bgp {

```



```

    export send-statics; #Allows advertising of a new route
group internal {
    type internal;
    local-address 192.168.1.1;
    neighbor 192.168.2.1;
    neighbor 192.168.4.1;
    neighbor 192.168.9.1;
    neighbor 192.168.6.1;
    neighbor 192.168.0.1;
}
}

ospf { #IGP enabled
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-0/1/0.0;
        interface fe-0/1/2.0;
        interface lo0.0 {
            passive; #Ensures protocols do not run over this interface
        }
    }
}

}

policy-options { #Load balancing policy
    policy-statement lbpp {
        then {
            load-balance per-packet;
        }
    }

    policy-statement send-statics { #Static route policy
        term statics {
            from {
                route-filter 100.100.1.0/24 exact;
            }
            then accept;
        }
    }
}
}

```

Sample Output 3

The following configuration output is for transit router **R2**.


```

user@R2> show configuration | no-more
[...Output truncated...]
interfaces {
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.0.24.1/30;
      }
      family mpls; #MPLS enabled on relevant interfaces
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.0.29.1/30;
      }
      family mpls;
    }
  }
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.12.14/30;
      }
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.144/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.2.1/32;
      }
    }
  }
}
routing-options {
  static {

```



```

[...Output truncated...]
router-id 192.168.2.1; #Manually configured RID
autonomous-system 65432; #Full mesh IBGP
}
}
protocols {
  rsvp {
    interface so-0/0/1.0;
    interface fe-0/1/0.0;
    interface so-0/0/2.0;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    interface fe-0/1/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group internal {
      type internal;
      local-address 192.168.2.1;
      neighbor 192.168.1.1;
      neighbor 192.168.4.1;
      neighbor 192.168.9.1;
      neighbor 192.168.6.1;
      neighbor 192.168.0.1;
    }
  }
  ospf { #IGP enabled
    traffic-engineering;
    area 0.0.0.0 {
      interface fe-0/1/0.0;
      interface so-0/0/1.0;
      interface so-0/0/2.0;
      interface lo0.0 {
        passive; #Ensures protocols do not run over this interface
      }
    }
  }
}

```



```

    }
}

```

Sample Output 4

The following configuration output is for transit router **R4**.

```

user@R4> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/30;
            }
            family mpls; # MPLS enabled on relevant interfaces
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.0.49.1/30;
            }
            family mpls;
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.146/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.4.1/32;
            }
        }
    }
}
routing-options {
    static {

```



```

[...Output truncated...]
router-id 192.168.4.1; #Manually configured RID
autonomous-system 65432; #Full mesh IBGP
}
protocols {
  rsvp {
    interface so-0/0/1.0;
    interface so-0/0/3.0;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    interface so-0/0/1.0;
    interface so-0/0/3.0;
    interface fxp0.0 {
      disable;
    }
  }
  bgp {
    group internal {
      type internal;
      local-address 192.168.4.1;
      neighbor 192.168.1.1;
      neighbor 192.168.2.1;
      neighbor 192.168.9.1;
      neighbor 192.168.6.1;
      neighbor 192.168.0.1;
    }
  }
  ospf { #IGP enabled
    traffic-engineering;
    area 0.0.0.0 {
      interface so-0/0/1.0;
      interface so-0/0/3.0;
      interface lo0.0 {
        passive; #Ensures protocols do not run over this interface
      }
    }
  }
}

```


Sample Output 5

The following configuration output is for transit router **R9**.

```
user@R9> show configuration | no-more
[...Output truncated...]
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.0.29.2/30;
      }
      family mpls; #MPLS enabled on relevant interfaces
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.0.49.2/30;
      }
      family mpls;
    }
  }
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.90.13/30;
      }
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.69.206/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.9.1/32;
      }
    }
  }
}
```



```

}
routing-options {
    static {
        [...Output truncated...]

        router-id 192.168.9.1; #Manually configured RID
        autonomous-system 65432; #Full mesh IBGP
    }
}
protocols {
    rsvp {
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fe-0/1/0.0;
        interface fxp0.0 {
            disable;
        }
    }
}
mpls {
    interface so-0/0/2.0;
    interface so-0/0/3.0;
    interface fe-0/1/0.0;
    interface fxp0.0 {
        disable;
    }
}
}
bgp {
    group internal {
        type internal;
        local-address 192.168.9.1;
        neighbor 192.168.1.1;
        neighbor 192.168.2.1;
        neighbor 192.168.4.1;
        neighbor 192.168.0.1;
        neighbor 192.168.6.1;
    }
}
}
ospf { #IGP enabled
    traffic-engineering;
    area 0.0.0.0 {
        interface so-0/0/2.0;
        interface so-0/0/3.0;
        interface fe-0/1/0.0;
        interface lo0.0 {
            passive; #Ensures protocols do not run over this interface
        }
    }
}

```



```

    }
  }
}

```

Sample Output 6

The following configuration output is for egress router **R0**.

```

user@R0> show configuration | no-more
[...Output truncated...]
interfaces {
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.90.14/30;
      }
      family mpls; #MPLS enabled on relevant interfaces
    }
  }
  fe-1/3/0 {
    unit 0 {
      family inet {
        address 10.10.11.1/24;
      }
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.69.207/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.0.1/32;
      }
    }
  }
}
routing-options {
  static {
    [...Output truncated...]
  }
}

```



```

        route 100.100.10.0/24 reject; #Static route for send-statics policy
    }
    router-id 192.168.0.1; #Manually configured RID
    autonomous-system 65432; #Full mesh IBGP
}
protocols {
    rsvp {
        interface fe-0/1/0.0;
        interface fe-1/3/0.0;
        interface fxp0.0 {
            disable;
        }
    }
}
mpls {
    label-switched-path r0-r6 {
        to 192.168.6.1;
    }
    interface fe-0/1/0.0;
    interface fe-1/3/0.0;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group internal {
        type internal;
        local-address 192.168.0.1;
        export send-statics; #Allows advertising of a new route
        neighbor 192.168.9.1;
        neighbor 192.168.6.1;
        neighbor 192.168.1.1;
        neighbor 192.168.2.1;
        neighbor 192.168.4.1;
    }
}
ospf { #IGP enabled
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-0/1/0.0;
        interface fe-1/3/0.0;
        interface lo0.0 {
            passive; #Ensures protocols do not run over this interface
        }
    }
}

```



```

    }
  }
}
policy-options {
  policy-statement send-statics {
    term statics {
      from {
        route-filter 100.100.10.0/24 exact;
      }
      then accept;
    }
  }
}
}

```

Meaning

Sample Outputs 1 through 6 show the base interfaces, routing options, protocols, and policy options configurations for all six routers in the example network illustrated in [“Example: Load-Balanced MPLS Network” on page 2101](#).

All routers in the network have MPLS, RSVP, and BGP enabled. OSPF is configured as the IGP, and relevant interfaces have basic IP information and MPLS support.

In addition, all routers have the router ID (RID) configured manually at the **[edit routing-options]** hierarchy level to avoid duplicate RID problems. The **passive** statement is included in the OSPF configuration to ensure that protocols are not run over the loopback (**lo0**) interface and that the loopback (**lo0**) interface is advertised correctly throughout the network.

Sample Outputs 1, 3, 4, and 5 for **R6**, **R2**, **R4**, and **R9** show the base configuration for transit label-switched routers. The base configuration includes all interfaces enabled for MPLS, the RID manually configured, and the relevant protocols (RSVP, MPLS, BGP, and OSPF).

Sample Output 2 from ingress router **R1** shows the base configuration plus four LSPs (**lsp1** through **lsp4**) configured to **R0**. The four LSPs are configured with different primary paths that specify a loose hop through **R4** for **lsp1** and **lsp4**, and through **R2** for **lsp2** and **lsp3**.

To create traffic, **R1** has a static route (**100.100.1.0/24**) configured at the **[edit routing-options static route]** hierarchy level. The prefix is included in the send-statics policy at the **[edit policy-options send statics]** hierarchy level so the routes can become BGP routes.

In addition, on the ingress router **R1**, load balancing is configured using the **per-packet** option, and the policy is exported at the **[edit routing-options forwarding-table]** hierarchy level.

Sample Output 6 from egress router **R0** shows one LSP (**r0-r6**) to **R6** used to create bidirectional traffic. OSPF requires bidirectional LSP reachability before it will advertise the LSP into the IGP. Although the

LSP is advertised into the IGP, no hello messages or routing updates occur over the LSP—only user traffic is sent over the LSP. The router uses its local copy of the IGP database to verify bidirectional reachability.

In addition, **R0** has a static route (**100.100.10.0/24**) configured at the **[edit routing-options static route]** hierarchy level. The prefix is included in the send-statics policy at the **[edit policy-options send statics]** hierarchy level so the routes can become BGP routes.

Traffic Flows Before Load Balancing

Purpose

The following sample output illustrates the details to look for when you issue different **show** commands to check if traffic is balanced. The following output is before load balancing is configured and is taken from transit router R2 in the network shown in [“Load-Balancing Network Topology” on page 2101](#).

Action

To check the distribution of traffic across interfaces and LSPs, use the following CLI operational mode commands:

```
user@host> show route | find mpls
user@host> monitor interface traffic
user@host> show mpls lsp statistics
```

Sample Output 1

```
user@R2> show route | find mpls

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 1d 00:12:08, metric 1
            Receive
1          *[MPLS/0] 1d 00:12:08, metric 1
            Receive
2          *[MPLS/0] 1d 00:12:08, metric 1
            Receive
100112     *[RSVP/7] 13:10:36, metric 1
            > via so-0/0/1.0, label-switched-path lsp1
100128     *[RSVP/7] 13:01:08, metric 1
            > via so-0/0/1.0, label-switched-path lsp4
100144     *[RSVP/7] 00:26:49, metric 1
            > to 10.0.12.13 via fe-0/1/0.0, label-switched-path r0-r6
```



```

100160      *[RSVP/7] 00:23:25, metric 1
             > via so-0/0/2.0, label-switched-path lsp2
100176      *[RSVP/7] 00:23:25, metric 1
             > via so-0/0/2.0, label-switched-path lsp3

```

Sample Output 2

```
user@R2> monitor interface traffic
```

```

R2                               Seconds: 89                               Time: 14:33:09

Interface    Link  Input packets      (pps)      Output packets      (pps)
so-0/0/0      Up           0           (0)           0           (0)
so-0/0/1 Up      90      (1)      91      (1)
so-0/0/2 Up     118     (1)    100122   (0)
so-0/0/3      Up           0           (0)           0           (0)
fe-0/1/0      Up    100119           (0)        115           (0)
fe-0/1/1      Up           0           (0)           0           (0)
fe-0/1/2      Up           0           (0)           0           (0)
fe-0/1/3      Up           0           (0)           0           (0)
[...Output truncated...]

```

Sample Output 3

```
user@R2> show mpls lsp statistics
```

```
Ingress LSP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

```
Egress LSP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

```
Transit LSP: 5 sessions
```

To	From	State	Packets	Bytes	LSPname
192.168.0.1	192.168.6.1	Up	0	0	lsp2
192.168.0.1	192.168.6.1	Up	112026	22853304	lsp1
192.168.0.1	192.168.6.1	Up	0	0	lsp3
192.168.0.1	192.168.6.1	Up	0	0	lsp4
192.168.6.1	192.168.0.1	Up	0	0	r0-r6
Total 5 displayed, Up 5, Down 0					

Meaning

Sample Outputs 1 through 3 from transit router **R2** show that traffic is not balanced across LSPs or interfaces.

Sample Output 1 for the **show route** command shows that all LSPs have the same metric (1) to the destination, even though they are traversing different interfaces. **lsp1** and **lsp4** are using **so-0/0/1**, while **lsp2** and **lsp3** are using **so-0/0/2**.

Sample Output 2 for the **monitor interface traffic** command shows that traffic is not evenly balanced across interfaces **so-0/0/1** and **so-0/0/2**. Almost all traffic is going out **so-0/0/2**.

Sample Output 3 for the **show mpls lsp statistics** command shows that traffic across LSPs is not balanced. All traffic is going over **lsp1**.

Related Topics

For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- *Junos User Guide*
- *Junos MPLS Applications Configuration Guide*
- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering*. White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

The Junos OS uses the load-balancing function across different protocols and features. For information about other types of load balancing, see the following:

- Option: Optimizing VPLS Traffic Flows, *Junos User Guide*
- Protocol-Independent Load Balancing for Layer 3 VPNs, *Junos VPNs Configuration Guide*
- Load Balancing Among Multiple Monitoring Interfaces, *Junos Services Interfaces Configuration Guide*

Verify the Operation of Uneven Bandwidth Load Balancing

Purpose

When a router is performing unequal-cost load balancing between LSPs paths, the **show route detail** command displays a balance field associated with each next hop being used.

Action

To verify that an RSVP LSP is unevenly load-balanced, use the following Junos OS CLI operational mode commands:


```
user@host> show route protocol rsvp detail
user@host> show mpls lsp statistics
```

Sample Output

```
user@R1> show route protocol rsvp detail

inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
10.0.90.14/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 7
          Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 10%
            Label-switched-path lsp1
            Label operation: Push 100768
            Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 20%
              Label-switched-path lsp2
              Label operation: Push 100736
              Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 30%,
selected
              Label-switched-path lsp3
              Label operation: Push 100752
              Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1  balance 40%
                Label-switched-path lsp4
                Label operation: Push 100784
                State: <Active Int>
                Local AS: 65432
                Age: 8:03      Metric: 4
                Task: RSVP
                Announcement bits (2): 0-KRT 4-Resolve tree 1
                AS path: I
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
192.168.0.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 7
          Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 10%
            Label-switched-path lsp1
            Label operation: Push 100768
            Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 20%
```



```

Label-switched-path lsp2
Label operation: Push 100736
Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 30%
Label-switched-path lsp3
Label operation: Push 100752
Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 40%, selected

```

```

Label-switched-path lsp4
Label operation: Push 100784
State: <Active Int>
Local AS: 65432
Age: 8:03      Metric: 4
Task: RSVP
Announcement bits (1): 1-Resolve tree 1
AS path: I

```

```
user@R1> show mpls lsp statistics
```

```
Ingress LSP: 4 sessions
```

To	From	State	Packets	Bytes	LSPname
192.168.0.1	192.168.1.1	Up	10067	845628	lsp1
192.168.0.1	192.168.1.1	Up	20026	1682184	lsp2
192.168.0.1	192.168.1.1	Up	29796	2502864	lsp3
192.168.0.1	192.168.1.1	Up	40111	3369324	lsp4

```
Total 4 displayed, Up 4, Down 0
```

```
Egress LSP: 1 sessions
```

To	From	State	Packets	Bytes	LSPname
192.168.1.1	192.168.0.1	Up	NA	NA	r0-r1

```
Total 1 displayed, Up 1, Down 0
```

```
Transit LSP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

Meaning

The sample output from ingress router **R1** shows that traffic is distributed according to the LSP bandwidth configuration, as indicated by the **Balance: xx%** field. For example, **lsp1** has 10 Mbps of bandwidth configured, as reflected in the **Balance: 10%** field.

Checklist for Collecting Crash Data

Problem

Description: [Table 60 on page 2122](#) provides links commands for collection crash data.

Solution

Table 60: Checklist for Collecting Crash Data

Tasks	Command or Action
“Understand Crash Data Collection” on page 2123	
“Collect Crash Data for a Routing Engine Kernel” on page 2124	
1. Check the Routing Engine Core Files on page 2124	<code>file list detail /var/crash</code>
2. Clear the NVRAM Contents on page 2147	
a. List the Core Files on page 2125	<pre>start shell su root password cd /var/crash ls -l</pre>
b. Compress the vmcore File on page 2126	<pre>gzip vmcore.number</pre> <p>To unzip the vmcore file:</p> <pre>gzip -d vmcore.number.gz</pre>
c. Log Software Version Information on page 2131	<code>show version</code>
d. Open a Case with JTAC on page 2128	Send an e-mail to support@juniper.net .
“Collect Crash Data for Routing Engine Daemons” on page 2132	
1. Check for Daemon Core Files on page 2132	<code>file list detail /var/tmp</code>
2. Collect and Send Routing Engine Crash Data to JTAC on page 2137	
a. List the Daemon Core Files on page 2134	<pre>start shell su root password cd /var/tmp ls -l</pre>
b. Compress the Daemon Core Files on page 2135	<code>gzip daemon-executable-name.core.number</code>
c. Log Software Version Information on page 2131	<code>show version</code>

Table 60: Checklist for Collecting Crash Data (*continued*)

Tasks	Command or Action
d. Open a Case with JTAC on page 2128	Send an e-mail to support@juniper.net .
“Collect Crash Data for the Packet Forwarding Engine Microkernel” on page 2141	
1. Display the Crash Stack Traceback and Registration Information on page 2143	<pre> start shell su root password vty component-executable-name show nvram show syslog messages </pre>
2. Clear the NVRAM Contents on page 2147	<pre> start shell su root password vty component-executable-name clear nvram </pre>
3. Check Packet Forwarding Engine Microkernel Core Files on page 2148	<code>file list detail /var/crash</code>
4. Collect and Send Routing Engine Crash Data to JTAC on page 2137	
a. List the Core Files Generated by the Crash on page 2148	<pre> start shell su root password cd /var/crash ls -l </pre>
b. Compress the Core Files on page 2149	<code>gzip filename</code>
c. Log Software Version Information on page 2131	<code>show version</code>
d. Open a Case with JTAC on page 2128	Send an e-mail to support@juniper.net .

Understand Crash Data Collection

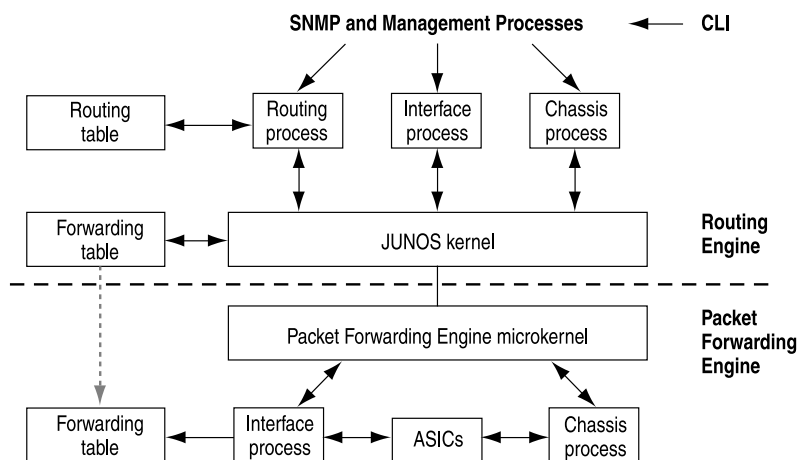
A crash can occur in one of three areas in the Junos OS architecture (see [Figure 162 on page 2124](#)):

- Routing Engine kernel

- Routing Engine daemons (processes)
- Packet Forwarding Engine microkernel

Figure 162: Three Areas Where a Software Crash Can Occur

Software Architecture



Collect Crash Data for a Routing Engine Kernel

Purpose

When a Routing Engine kernel crashes, the Routing Engine automatically reboots. By default, the Juniper Networks router does not attempt to dump a core if the Routing Engine kernel crashes. As a result, there is no crash data on the router to help investigate the crash. In addition, the system log messages are similar to those generated when the router is powered down and restarted, so you cannot tell if the Routing Engine restart was caused by a kernel crash or a normal power restart.

To collect crash data for a Routing Engine kernel crash, follow these steps:

1. [Check the Routing Engine Core Files | 2124](#)
2. [List the Core Files | 2125](#)
3. [Compress the vmcore File | 2126](#)
4. [Log Software Version Information | 2127](#)
5. [Open a Case with JTAC | 2128](#)

Check the Routing Engine Core Files

Purpose

If you observe disruption to the Routing Engine kernel, check the **/var/crash** directory for any core files created around the time of the crash.

Action

To check the **/var/crash** directory, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> file list detail /var/crash
```

Sample Output

```
user@host> file list detail /var/crash
```

```
total 1577912
drwxr-x---  2 root  wheel      512 Sep  9 11:59 ./
drwxr-xr-x 22 root  wheel      512 Oct 29 2001 ../
-rw-r--r--  1 root  wheel        2 Jul 20 01:11 bounds
-rw-r--r--  1 root  wheel 2166913 Jul 20 01:11 kernel.0
-rw-r--r--  1 root  wheel        5 Feb 15 2002 minfree
-rw-----  1 root  wheel 805306368 Jul 20 01:11 vmcore.0
```

Meaning

The sample output lists the contents of the **/var/crash/** directory. Check the date and timestamp for any kernel core files created around the time of the crash. In the example above, two core files are listed: **kernel.0** and **vmcore.0**.

List the Core Files

Purpose

To list the core files, follow these steps:

Action

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and the prompt is % instead of >, for example:


```
% su
Password: ****
```

3. Change the directory to **/var/crash** and type **ls -l**, for example:

```
root@host% cd /var/crash
root@host% ls -l
```

4. Look for any core files created around the time of the crash.

Sample Output

user@host> **start shell**

```
% su
Password: ****
root@host% cd /var/crash
root@host% ls -l
total 1577908
-rw-r--r-- 1 root wheel          2 Jul 20 01:11 bounds
-rw-r--r-- 1 root wheel 2166913 Jul 20 01:11 kernel.0
-rw-r--r-- 1 root wheel          5 Feb 15 2002 minfree
-rw----- 1 root wheel 805306368 Jul 20 01:11 vmcore.0
```

Meaning

The sample output lists the contents of the **/var/crash** directory and shows the current core files **kernel.0** and **vmcore.0**.

Compress the vmcore File

Purpose

The **gzip** compression utility is used to compress files. Compress the **vmcore** file if it is larger than 50 MB. Files created using the **gzip** command end with the file extension **.gz**.

NOTE: Use lowercase for the **gzip** command when you are in the shell.

Action

To compress the **vmcore** file with gzip, use the following command from the shell:

```
root@host% gzip vmcore .number
```

To unzip the **vmcore** file with gzip, use the following command from the shell:

```
root@host% gzip -d vmcore .number.gz
```

Meaning

The contents of the **vmcore** file are compressed into a single compressed file named **vmcore.number.gz**. The **gzip** command preserves the mode, ownership, and timestamps of files when compressing or decompressing them.

Log Software Version Information

Purpose

To log the Junos OS version information.

Action

To log the Junos OS version information, use the following Junos OS CLI operational mode command:

```
user@host> show version
```

Sample Output

```
user@host> show version
```

```
Hostname: host
Model: m10
JUNOS Base OS boot [5.0R5]
JUNOS Base OS Software Suite [5.0R5]
JUNOS Kernel Software Suite [5.0R5]
JUNOS Routing Software Suite [5.0R5]
JUNOS Packet Forwarding Engine Support [5.0R5]
JUNOS Crypto Software Suite [5.0R5]
JUNOS Online Documentation [5.0R5]
KERNEL 5.0R5 #0 built by builder on 2002-03-02 05:10:28 UTC
MGD release 5.0R5 built by builder on 2002-03-02 04:45:32 UTC
CLI release 5.0R5 built by builder on 2002-03-02 04:44:22 UTC
CHASSISD release 5.0R5 built by builder on 2002-03-02 04:43:37 UTC
DCD release 5.0R5 built by builder on 2002-03-02 04:42:47 UTC
```



```

RPD release 5.0R5 built by builder on 2002-03-02 04:46:17 UTC
SNMPD release 5.0R5 built by builder on 2002-03-02 04:52:26 UTC
MIB2D release 5.0R5 built by builder on 2002-03-02 04:45:37 UTC
APSD release 5.0R5 built by builder on 2002-03-02 04:43:31 UTC
VRRPD release 5.0R5 built by builder on 2002-03-02 04:52:34 UTC
ALARMD release 5.0R5 built by builder on 2002-03-02 04:43:24 UTC
PFED release 5.0R5 built by builder on 2002-03-02 04:46:06 UTC
CRAFTD release 5.0R5 built by builder on 2002-03-02 04:44:30 UTC
SAMPLED release 5.0R5 built by builder on 2002-03-02 04:52:20 UTC
ILMID release 5.0R5 built by builder on 2002-03-02 04:45:21 UTC
BPRELAYD release 5.0R5 built by builder on 2002-03-02 04:42:41 UTC
RMOPD release 5.0R5 built by builder on 2002-03-02 04:46:11 UTC
jkernel-dd release 5.0R5 built by builder on 2002-03-02 04:41:07 UTC
jroute-dd release 5.0R5 built by builder on 2002-03-02 04:41:21 UTC
jdocs-dd release 5.0R5 built by builder on 2002-03-02 04:39:11 UTC

```

Meaning

The sample output shows the hostname, router model, and the different Junos OS packages, processes, and documents.

Open a Case with JTAC

Problem

Description: To open a case with JTAC, call one of the Support phone numbers or create a case via [Case Manager](#).

Check the Routing Engine Core Files

Purpose

If you observe disruption to the Routing Engine kernel, check the **/var/crash** directory for any core files created around the time of the crash.

Action

To check the **/var/crash** directory, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> file list detail /var/crash
```


Sample Output

user@host> **file list detail /var/crash**

```
total 1577912
drwxr-x---  2 root  wheel      512 Sep  9 11:59 ./
drwxr-xr-x 22 root  wheel      512 Oct 29  2001 ../
-rw-r--r--  1 root  wheel        2 Jul 20 01:11 bounds
-rw-r--r--  1 root  wheel 2166913 Jul 20 01:11 kernel.0
-rw-r--r--  1 root  wheel        5 Feb 15  2002 minfree
-rw-----  1 root  wheel 805306368 Jul 20 01:11 vmcore.0
```

Meaning

The sample output lists the contents of the **/var/crash/** directory. Check the date and timestamp for any kernel core files created around the time of the crash. In the example above, two core files are listed: **kernel.0** and **vmcore.0**.

List the Core Files

Purpose

To list the core files, follow these steps:

Action

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and the prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Change the directory to **/var/crash** and type **ls -l**, for example:

```
root@host% cd /var/crash
```



```
root@host% ls -l
```

4. Look for any core files created around the time of the crash.

Sample Output

```
user@host> start shell
```

```
% su
Password: ****
root@host% cd /var/crash
root@host% ls -l
total 1577908
-rw-r--r--  1 root  wheel           2 Jul 20 01:11 bounds
-rw-r--r--  1 root  wheel    2166913 Jul 20 01:11 kernel.0
-rw-r--r--  1 root  wheel           5 Feb 15 2002 minfree
-rw-----  1 root  wheel   805306368 Jul 20 01:11 vmcore.0
```

Meaning

The sample output lists the contents of the **/var/crash** directory and shows the current core files **kernel.0** and **vmcore.0**.

Compress the vmcore File

Purpose

The gzip compression utility is used to compress files. Compress the **vmcore** file if it is larger than 50 MB. Files created using the **gzip** command end with the file extension **.gz**.

NOTE: Use lowercase for the **gzip** command when you are in the shell.

Action

To compress the **vmcore** file with gzip, use the following command from the shell:

```
root@host% gzip vmcore .number
```


To unzip the **vmcore** file with **gzip**, use the following command from the shell:

```
root@host% gzip -d vmcore .number.gz
```

Meaning

The contents of the **vmcore** file are compressed into a single compressed file named **vmcore.number.gz**. The **gzip** command preserves the mode, ownership, and timestamps of files when compressing or decompressing them.

Log Software Version Information

Purpose

To log the JUNOS software version information.

Action

To log the JUNOS software version information, use the following JUNOS CLI operational mode command:

```
user@host> show version
```

Sample Output

```
user@host> show version
```

```
Hostname: host
Model: m10
JUNOS Base OS boot [5.0R5]
JUNOS Base OS Software Suite [5.0R5]
JUNOS Kernel Software Suite [5.0R5]
JUNOS Routing Software Suite [5.0R5]
JUNOS Packet Forwarding Engine Support [5.0R5]
JUNOS Crypto Software Suite [5.0R5]
JUNOS Online Documentation [5.0R5]
KERNEL 5.0R5 #0 built by builder on 2002-03-02 05:10:28 UTC
MGD release 5.0R5 built by builder on 2002-03-02 04:45:32 UTC
CLI release 5.0R5 built by builder on 2002-03-02 04:44:22 UTC
CHASSISD release 5.0R5 built by builder on 2002-03-02 04:43:37 UTC
DCD release 5.0R5 built by builder on 2002-03-02 04:42:47 UTC
RPD release 5.0R5 built by builder on 2002-03-02 04:46:17 UTC
SNMPD release 5.0R5 built by builder on 2002-03-02 04:52:26 UTC
MIB2D release 5.0R5 built by builder on 2002-03-02 04:45:37 UTC
```



```

APSD release 5.0R5 built by builder on 2002-03-02 04:43:31 UTC
VRRPD release 5.0R5 built by builder on 2002-03-02 04:52:34 UTC
ALARMD release 5.0R5 built by builder on 2002-03-02 04:43:24 UTC
PFED release 5.0R5 built by builder on 2002-03-02 04:46:06 UTC
CRAFTD release 5.0R5 built by builder on 2002-03-02 04:44:30 UTC
SAMPLED release 5.0R5 built by builder on 2002-03-02 04:52:20 UTC
ILMID release 5.0R5 built by builder on 2002-03-02 04:45:21 UTC
BPRELAYD release 5.0R5 built by builder on 2002-03-02 04:42:41 UTC
RMOPD release 5.0R5 built by builder on 2002-03-02 04:46:11 UTC
jkernel-dd release 5.0R5 built by builder on 2002-03-02 04:41:07 UTC
jroute-dd release 5.0R5 built by builder on 2002-03-02 04:41:21 UTC
jdocs-dd release 5.0R5 built by builder on 2002-03-02 04:39:11 UTC

```

Meaning

The sample output shows the hostname, router model, and the different JUNOS software packages, processes, and documents.

Open a Case with JTAC

Problem

Description: To open a case with JTAC, call one of the Support phone numbers or create a case via [Case Manager](#).

Collect Crash Data for Routing Engine Daemons

To collect crash data for Routing Engine daemons, follow these steps:

1. [Check for Daemon Core Files | 2132](#)
2. [List the Daemon Core Files | 2134](#)
3. [Compress the Daemon Core Files | 2135](#)
4. [Log Software Version Information | 2136](#)
5. [Open a Case with JTAC | 2137](#)

Check for Daemon Core Files

Purpose

If you observe disruption to routing protocol operation, system log operation, Simple Network Management Protocol (SNMP) operation, or other operations handled by Routing Engine daemons, check the **/var/tmp** directory for any daemon core files created around the time of the crash.

Action

To check the **/var/tmp** directory, use the following Junos OS CLI operational mode command:

```
user@host> file list detail /var/tmp
```

Sample Output

```
user@host> file list detail /var/tmp
```

```
total 1292622
drwxrwxrwt  3 root  field      512 Dec 31 06:48 ./
drwxr-xr-x 21 root  field      512 Mar  5 1999 ../
-rw-rw----  1 root  field 119713792 Nov 17 21:58 rpd.core.0
-rw-rw----  1 root  field 120782848 Nov 17 22:12 rpd.core.1
```

Meaning

The sample output lists the contents of the **/var/tmp/** directory. Look for any daemon core files created around the time of the crash. In the example above, two core files are listed: **rpd.core.0** and **rpd.core.1**.

[Table 61 on page 2133](#) lists the major Routing Engine daemons supported by the Junos OS.

Table 61: Major Routing Engine Daemons

Executable Name	Definition	Description
rpd	Routing protocol daemon	Provides routing protocol intelligence (Border Gateway Protocol [BGP], Intermediate System-to-Intermediate System [ISIS], Open Shortest Path First [OSPF], and so on).
dcd	Device control daemon	Manages all interface devices.
mgd	Management daemon	Provides user configuration access to the system. The CLI is a client of mgd .
snmpd	Simple Network Management Protocol daemon	Provides remote network management information to the network management system.

Table 61: Major Routing Engine Daemons (continued)

Executable Name	Definition	Description
chassisd	Chassis daemon	Monitors and manages Flexible PIC Concentrator (FPC) slots and other environmental components.
alarmd	Alarm daemon	Manages system alarm notifications.
apspd	Automatic protection switching daemon	Provides SONET Automatic Protection Switching (APS) functionality.
sampld	Traffic sampling daemon	Gathers traffic sampling information.
vrrpd	Virtual Router Redundancy Protocol daemon	Provides Virtual Router Redundancy Protocol (VRRP) functionality.
syslogd	System log daemon	Manages the router system logging operation.
mib2d	MIB2 daemon	Management Information Base (MIB) subagent for MIB2.

List the Daemon Core Files

Purpose

To list the daemon core files.

Action

To list the daemon core files, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and the prompt is **%** instead of **>**, for example:

```
% su
Password: ****
```

3. Change the directory to **/var/tmp** and type **ls -l**, for example:


```
root@host% cd /var/tmp
root@host% ls -l
```

4. Look for any daemon core files created around the time of the crash.

Sample Output

user@host> **start shell**

```
% su
Password: ****
root@host% cd /var/tmp
root@host% ls -l
total 1292618
-rw-rw----  1 root  field  119713792 Nov 17 21:58  rpd.core.0
-rw-rw----  1 root  field  120782848 Nov 17 22:12  rpd.core.1
```

Meaning

The sample output lists the contents of the **/var/tmp** directory and shows the current core file (**rpd.core.1**) and one previous core file (**rpd.core.0**) for the routing protocol daemon (**rpd**). For each daemon, you can have a total of five core files in the **/var/tmp** directory: the current core file and the four previous core files numbered 0 through 4 (from oldest to newest).

Compress the Daemon Core Files

Purpose

The **gzip** compression utility is used to compress the files if they are large. Files created using the **gzip** command end with the file extension **.gz**. Compress the core file if it is over 50 MB.

NOTE: Use lowercase for the **gzip** command when you are in the shell.

You only need to compress the daemon core files when the tarball file is not created.

Action

To compress the daemon core file with **gzip**, use the following command from the shell:


```
root@host% gzip daemon-executable-name.core.number
```

Sample Output

```
root@host% gzip rpd.core.0
gzip rpd.core.0
```

Meaning

The contents of the daemon core file are compressed into a single compressed file named **daemon.number.gz**. The **gzip** command preserves the mode, ownership, and timestamps of files when compressing or decompressing them.

Log Software Version Information

Purpose

To log the Junos OS version information.

Action

To log the Junos OS version information, use the following Junos OS CLI operational mode command:

```
user@host> show version
```

Sample Output

```
user@host> show version
```

```
Hostname: host
Model: m10
JUNOS Base OS boot [5.0R5]
JUNOS Base OS Software Suite [5.0R5]
JUNOS Kernel Software Suite [5.0R5]
JUNOS Routing Software Suite [5.0R5]
JUNOS Packet Forwarding Engine Support [5.0R5]
JUNOS Crypto Software Suite [5.0R5]
JUNOS Online Documentation [5.0R5]
KERNEL 5.0R5 #0 built by builder on 2002-03-02 05:10:28 UTC
MGD release 5.0R5 built by builder on 2002-03-02 04:45:32 UTC
CLI release 5.0R5 built by builder on 2002-03-02 04:44:22 UTC
CHASSISD release 5.0R5 built by builder on 2002-03-02 04:43:37 UTC
```



```

DCD release 5.0R5 built by builder on 2002-03-02 04:42:47 UTC
RPD release 5.0R5 built by builder on 2002-03-02 04:46:17 UTC
SNMPD release 5.0R5 built by builder on 2002-03-02 04:52:26 UTC
MIB2D release 5.0R5 built by builder on 2002-03-02 04:45:37 UTC
APSD release 5.0R5 built by builder on 2002-03-02 04:43:31 UTC
VRRPD release 5.0R5 built by builder on 2002-03-02 04:52:34 UTC
ALARMD release 5.0R5 built by builder on 2002-03-02 04:43:24 UTC
PFED release 5.0R5 built by builder on 2002-03-02 04:46:06 UTC
CRAFTD release 5.0R5 built by builder on 2002-03-02 04:44:30 UTC
SAMPLED release 5.0R5 built by builder on 2002-03-02 04:52:20 UTC
ILMID release 5.0R5 built by builder on 2002-03-02 04:45:21 UTC
BPRELAYD release 5.0R5 built by builder on 2002-03-02 04:42:41 UTC
RMOPD release 5.0R5 built by builder on 2002-03-02 04:46:11 UTC
jkernel-dd release 5.0R5 built by builder on 2002-03-02 04:41:07 UTC
jroute-dd release 5.0R5 built by builder on 2002-03-02 04:41:21 UTC
jdocs-dd release 5.0R5 built by builder on 2002-03-02 04:39:11 UTC

```

Meaning

The sample output shows the hostname, router model, and the different Junos OS packages, processes, and documents.

Open a Case with JTAC

Problem

Description: To open a case with JTAC, call one of the Support phone numbers or create a case via [Case Manager](#).

Collect and Send Routing Engine Crash Data to JTAC

Problem

Description: If a Routing Engine kernel crash occurs on your router, collect the following data for JTAC evaluation and instruction.

Solution

To collect and send Routing Engine crash data to JTAC, follow these steps:

1. [List the Core Files on page 2125](#)
2. [Compress the vmcore File on page 2126](#)
3. [Log Software Version Information on page 2131](#)
4. [Open a Case with JTAC on page 2128](#)

RELATED DOCUMENTATION

- [Understand Crash Data Collection | 2123](#)
- [Checklist for Collecting Crash Data | 2121](#)

Check for Daemon Core Files

Purpose

If you observe disruption to routing protocol operation, system log operation, Simple Network Management Protocol (SNMP) operation, or other operations handled by Routing Engine daemons, check the **/var/tmp** directory for any daemon core files created around the time of the crash.

Action

To check the **/var/tmp** directory, use the following Junos OS CLI operational mode command:

```
user@host> file list detail /var/tmp
```

Sample Output

```
user@host> file list detail /var/tmp
```

```
total 1292622
drwxrwxrwt   3 root  field      512 Dec 31 06:48 ./
drwxr-xr-x  21 root  field      512 Mar  5 1999 ../
-rw-rw----   1 root  field 119713792 Nov 17 21:58 rpd.core.0
-rw-rw----   1 root  field 120782848 Nov 17 22:12 rpd.core.1
```

Meaning

The sample output lists the contents of the **/var/tmp/** directory. Look for any daemon core files created around the time of the crash. In the example above, two core files are listed: **rpd.core.0** and **rpd.core.1**.

[Table 61 on page 2133](#) lists the major Routing Engine daemons supported by the Junos OS.

Table 62: Major Routing Engine Daemons

Executable Name	Definition	Description
rpdd	Routing protocol daemon	Provides routing protocol intelligence (Border Gateway Protocol [BGP], Intermediate System-to-Intermediate System [ISIS], Open Shortest Path First [OSPF], and so on).
dcd	Device control daemon	Manages all interface devices.
mgd	Management daemon	Provides user configuration access to the system. The CLI is a client of mgd .
snmpd	Simple Network Management Protocol daemon	Provides remote network management information to the network management system.
chassisd	Chassis daemon	Monitors and manages Flexible PIC Concentrator (FPC) slots and other environmental components.
alarmd	Alarm daemon	Manages system alarm notifications.
apsd	Automatic protection switching daemon	Provides SONET Automatic Protection Switching (APS) functionality.
sampled	Traffic sampling daemon	Gathers traffic sampling information.
vrripd	Virtual Router Redundancy Protocol daemon	Provides Virtual Router Redundancy Protocol (VRRP) functionality.
syslogd	System log daemon	Manages the router system logging operation.
mib2d	MIB2 daemon	Management Information Base (MIB) subagent for MIB2.

List the Daemon Core Files

Purpose

To list the daemon core files.

Action

To list the daemon core files, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and the prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Change the directory to **/var/tmp** and type **ls -l**, for example:

```
root@host% cd /var/tmp
root@host% ls -l
```

4. Look for any daemon core files created around the time of the crash.

Sample Output

```
user@host> start shell
```

```
% su
Password: ****
root@host% cd /var/tmp
root@host% ls -l
total 1292618
-rw-rw----  1 root  field  119713792 Nov 17 21:58  rpd.core.0
-rw-rw----  1 root  field  120782848 Nov 17 22:12  rpd.core.1
```

Meaning

The sample output lists the contents of the **/var/tmp** directory and shows the current core file (**rpd.core.1**) and one previous core file (**rpd.core.0**) for the routing protocol daemon (**rpd**). For each daemon, you can have a total of five core files in the **/var/tmp** directory: the current core file and the four previous core files numbered 0 through 4 (from oldest to newest).

Compress the Daemon Core Files

Purpose

The gzip compression utility is used to compress the files if they are large. Files created using the **gzip** command end with the file extension **.gz**. Compress the core file if it is over 50 MB.

NOTE: Use lowercase for the **gzip** command when you are in the shell.

You only need to compress the daemon core files when the tarball file is not created.

Action

To compress the daemon core file with gzip, use the following command from the shell:

```
root@host% gzip daemon-executable-name.core.number
```

Sample Output

```
root@host% gzip rpd.core.0
gzip rpd.core.0
```

Meaning

The contents of the daemon core file are compressed into a single compressed file named **daemon.number.gz**. The **gzip** command preserves the mode, ownership, and timestamps of files when compressing or decompressing them.

Collect Crash Data for the Packet Forwarding Engine Microkernel

Purpose

Each of the following Packet Forwarding Engine components of a Juniper Networks router runs a microkernel:

- Flexible PIC Concentrator (FPC) on M-series platforms except for the M5 and M10 Internet routers
- Gibson Flexible PIC Concentrator (GFPC) on T640 and T320 Internet routing nodes
- Switched Printed Mezzanine Board (SPMB) on T640 and T320 Internet routing nodes
- Forwarding Engine Board (FEB) on M5 and M10 Internet routers

- System Switching Board (SSB) on an M20 Internet router
- System Control Board (SCB) on an M40 Internet router
- Switching and Forwarding Module (SFM) on M160 and M40e Internet routers

When a crash occurs, crash stack traceback and registration information is placed into nonvolatile random access memory (NVRAM) on the different components. [Table 63 on page 2142](#) shows where the NVRAM is located for the components for each router.

Table 63: NVRAM Location on the Microkernel of the Packet Forwarding Engine Components

Router Type	NVRAM Location
M5 and M10	FEB
M20	SSB and crash stack traceback and register information for the FPC
M40	SCB and crash stack traceback and register information for the FPC
M40e	FPC SFM
M160	FPC SFM
T320	GFPC SPMB
T640	GFPC SPMB

To collect crash data for the Packet Forwarding Engine microkernel, follow these steps:

1. [Display the Crash Stack Traceback and Registration Information | 2143](#)
2. [Clear the NVRAM Contents | 2147](#)
3. [Check Packet Forwarding Engine Microkernel Core Files | 2148](#)
4. [List the Core Files Generated by the Crash | 2148](#)
5. [Compress the Core Files | 2149](#)
6. [Log Software Version Information | 2150](#)
7. [Open a Case with JTAC | 2151](#)

Display the Crash Stack Traceback and Registration Information

Purpose

To display the crash stack traceback and registration information.

Action

To display the crash stack traceback and registration information, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Establish a vty session to the appropriate component. Use the **vtty** command followed by the executable name for the component; for example, **scb**, **ssb0**, **ssb1**, **fpc0**, or **fpc1**:

```
root@host% vty sfm0
```

NOTE: For the M40e and M160 routers, you can also create a **cty** session to the components if the components are not online.

4. Type the **show nvram** command to view the NVRAM information.
5. Type the **show syslog messages** command to view the system log messages.

Sample Output 1

```
user@host> start shell
```

```
% su
Password: ****
root@host% vty sfm0
```


SFM platform (266Mhz PPC 603e processor, 64Mb memory, 512Kb flash)

SFM3(host vty)# **show nvram**

System NVRAM :

4080 available bytes, 4080 used, 0 free

Contents:

mpc106 machine check caused by error on the PCI Bus

mpc106 error detect register 1: 0x08, 2: 0x00

mpc106 error ack count = 0

mpc106 error address: 0x0a000000

mpc106 PCI bus error status register: 0x02

mpc106 was the PCI master

C/BE bits: I/O read [0b0010]

mpc106 error detection reg1: PCI cycle

mpc106 PCI status reg: parity error

System Exception: Vector/Code 0x00700, Signal 4

Event occurred at: Oct 26 13:32:40.952

Juniper Embedded Microkernel Version 4.2R1

Built by tlim on 2000-09-23 06:11:28 UTC

Copyright (C) 1998-2000, Juniper Networks, Inc.

All rights reserved.

Reason string: "Program Check"

Context: Thread (PFE Manager)

Registers:

R00: 0x06f5f81c R01: 0x06f5f9cc R02: 0x00003344 R03: 0x00000000

R04: 0x00008000 R05: 0x00000000 R06: 0x0010052c R07: 0x06f637e4

R08: 0x06f5f81c R09: 0x00169810 R10: 0x000000e8 R11: 0x00000001

R12: 0x00046cdf R13: 0xffffffff R14: 0xffffffff R15: 0xffffffff

R16: 0xffffffff R17: 0xffffffff R18: 0xffffffff R19: 0xffffffff

R20: 0xffffffff R21: 0xffffffff R22: 0xffffffff R23: 0xffffffff

R24: 0x00000003 R25: 0x00000000 R26: 0x00000001 R27: 0x0000fc78

R28: 0x00150000 R29: 0x0016c4b0 R30: 0x06f5eb7c R31: 0x97cb1d36

MSR: 0x0008b030 CTR: 0x000ac008 Link:0x06f5f81c SP: 0x06f5f9cc

CCR: 0x22200024 XER: 0x20000000 PC: 0x06f5f81c

DSISR: 0x00000000 DAR: 0xffffffff K_MSR: 0x00001030

Stack Traceback :

Frame 01: sp = 0x06f5f9cc, pc = 0x06f5f81c

Frame 02: sp = 0x06f5f9e4, pc = 0x000c7e28

Frame 03: sp = 0x06f5fa04, pc = 0x00026620

ROM NVRAM:

0 available bytes, 0 used, 0 free

SFM3(host vty)# **show syslog messages**

```
Oct 26 12:02:05 router tnp_sfm_2 PFEMAN: sent Resync request to Master
Oct 26 12:02:07 router tnp_sfm_3 CM(3): Slot 1: On-line
Oct 26 12:02:07 router tnp_sfm_3 CM(3): Slot 2: On-line
Oct 26 12:02:07 router tnp_sfm_3 CM(3): Slot 6: On-line
Oct 26 12:02:07 router tnp_sfm_3 PFEMAN: sent Resync request to Master
Oct 26 12:05:58 router tnp_sfm_3 mpc106 machine check caused by error on the
PCI Bu
s
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error detect register 1: 0x08,
2: 0x00
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error ack count = 0
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error address: 0x0a000000
Oct 26 12:05:58 router tnp_sfm_3 mpc106 PCI bus error status register: 0x02
Oct 26 12:05:58 router tnp_sfm_3 mpc106 was the PCI master
Oct 26 12:05:58 router tnp_sfm_3 C/BE bits: I/O read [0b0010]
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error detection reg1: PCI cycle
Oct 26 12:05:58 router tnp_sfm_3 mpc106 PCI status reg: parity error
Oct 26 12:05:58 router tnp_sfm_3 ^B
Oct 26 12:05:58 router tnp_sfm_3 last message repeated 7 times
Oct 26 12:05:58 router tnp_sfm_3 Registers:
Oct 26 12:05:58 router tnp_sfm_3 R00: 0x06f5f81c R01: 0x06f5f9cc
R02: 0x00003344 R0
3: 0x00000000
Oct 26 12:05:58 router tnp_sfm_3 R04: 0x00008000 R05: 0x00000000
R06: 0x0010052c R0
7: 0x06f637e4
Oct 26 12:05:58 router tnp_sfm_3 R08: 0x06f5f81c R09: 0x00169810
R10: 0x000003b4 R1
1: 0x00000001
Oct 26 12:05:58 router tnp_sfm_3 R12: 0x00017b97 R13: 0xffffffff
R14: 0xffffffff R1
5: 0xffffffff
Oct 26 12:05:58 router tnp_sfm_3 R16: 0xffffffff R17: 0xffffffff
R18: 0xffffffff R1
9: 0xffffffff
Oct 26 12:05:58 router tnp_sfm_3 R20: 0xffffffff R21: 0xffffffff
R22: 0xffffffff R2
3: 0xffffffff
```



```

Oct 26 12:05:58 router tnp_sfm_3 R24: 0x00000003 R25: 0x00000000
R26: 0x00000001 R2
7: 0x0000fc78
Oct 26 12:05:58 router tnp_sfm_3 R28: 0x00150000 R29: 0x0016c4b0
R30: 0x06f5eb7c R3
1: 0x97c9c35e
Oct 26 12:05:58 router tnp_sfm_3 MSR: 0x0008b030 CTR: 0x000ac008
Link:0x06f5f81c SP
: 0x06f5f9cc
Oct 26 12:05:58 router tnp_sfm_3 CCR: 0x22200024 XER: 0x20000000
PC: 0x06f5f81c
Oct 26 12:05:58 router tnp_sfm_3 DSISR: 0x00000000 DAR: 0xffffffff
K_MSR: 0x0000103
0

```

Sample Output

The following sample output is another example of displaying the crash stack traceback and registration information:

```

root@host% vty fpc1
FPC160 platform (PPC 603e processor, 32Mb memory, 512Kb flash)

FPC1(host vty)# show nvram
System NVRAM :
  4080 available bytes, 4080 used, 0 free
  Contents:
0000000 R06: 0x0000005c R07: 0x850400d0
R08: 0x00000000 R09: 0x00000020 R10: 0x00000000 R11: 0x00000129
R12: 0x00000000 R13: 0x00000000 R14: 0x4005009a R15: 0x20000260
R16: 0xc8828784 R17: 0x84212800 R18: 0xc0004c61 R19: 0x80005900
R20: 0x80206000 R21: 0x84000304 R22: 0xd0410180 R23: 0x8c2005ac
R24: 0x00000003 R25: 0x00000000 R26: 0x00000001 R27: 0x0000fc48
R28: 0x001d0000 R29: 0x00000001 R30: 0x00136bb8 R31: 0x00000000
MSR: 0x0000b030 CTR: 0x001331e0 Link:0x000308c8 SP: 0x01baba34
CCR: 0x42200020 XER: 0x00000000 PC: 0x000308cc
DSISR: 0x00000000 DAR: 0xffffffff K_MSR: 0x00001030

Stack Traceback:
Frame 01: sp = 0x01baba34, pc = 0x000308c8
Frame 02: sp = 0x01babac4, pc = 0x0002647c

```



```

Frame 03: sp = 0x01babad4, pc = 0x00026590
Frame 04: sp = 0x01babadc, pc = 0x00106fcc
Frame 05: sp = 0x01babafc, pc = 0x00026620
ROM NVRAM:
  0 available bytes, 0 used, 0 free

FPC1(host vty)# show syslog messages
[0+00:00:00.780 LOG: Info] Version 4.0R5 by tlim on 2000-08-10 04:45:54 UTC
[0+00:00:00.780 LOG: Info] On-board NVRAM contains diagnostic information.
[0+00:00:03.175 LOG: Info] PFEMAN: Established connection to Master
[Jan 30 21:53:05.804 LOG: Info] SNTPD: Initial time of day set.

```

Meaning

Sample output 1 and 2 show the stack trace from the microkernel crash. Save the output from the **show nvram** and **show syslog** commands so that you can send them to JTAC when you open a case.

Clear the NVRAM Contents

Purpose

Currently the storage area for the logs on the NVRAM is limited to 4 KB. You need to delete old NVRAM logs to make room for new ones.

Action

To clear the content of the NVRAM after you have captured the necessary information, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and the prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Establish a vty session to the appropriate component. Use the **vtty** command followed by the abbreviation for the component, for example:

```
root@host% vtty sfm0
SFM3(host vty)#
```



```
FPC1(host vty)#
```

4. Type the **clear nvram** command, for example:

```
SFM3(host vty)# clear nvram
FPC1(host vty)# clear nvram
```

Check Packet Forwarding Engine Microkernel Core Files

Purpose

If you observe disruption to the Packet Forwarding Engine microkernel, check the **/var/crash** directory for any core files created around the time of the crash.

Action

To check the **/var/crash** directory, use the following Junos OS CLI operational mode command:

```
user@host> file list detail /var/crash
```

Sample Output

```
user@host> file list detail /var/crash
```

```
var/crash:
total 456630
-rw-r--r--  1 root  wheel    6814720 Dec 18 08:03 core-FPC4.100111808032
-rw-r--r--  1 root  wheel    65613824 Dec 10 04:58 core-SCB.100111004570
-rw-r--r--  1 root  wheel    65613824 Dec 19 00:23 core-SCB.100111900221
-rw-r--r--  1 root  wheel    65545216 Feb  9 20:46 core-SCB.101010920452
```

Meaning

The sample output lists the contents of the **/var/crash/** directory. Check the date and timestamp for any core files created around the time of the crash. In the example above, four core files are listed.

List the Core Files Generated by the Crash

Purpose

To list the core files generated by the crash.

Action

To list the core files, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and the prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Change the directory to **/var/crash** and type **ls -l**, for example:

```
root@host% cd /var/crash
root@host% ls -l
```

4. Look for any core files created around the time of the crash.

Sample Output

```
user@host> start shell
```

```
% su
Password: ****
root@host% cd /var/crash
root@host% ls -l
total 456630
-rw-r--r--  1 root  wheel    6814720 Dec 18 08:03 core-FPC4.100111808032
-rw-r--r--  1 root  wheel    65613824 Dec 10 04:58 core-SCB.100111004570
-rw-r--r--  1 root  wheel    65613824 Dec 19 00:23 core-SCB.100111900221
-rw-r--r--  1 root  wheel    65545216 Feb  9 20:46 core-SCB.101010920452
```

Meaning

The sample output shows the current core files for the different components on the router; for example, **core-FPC4.100111808032** and **core-SCB.100111004570**.

Compress the Core Files

Purpose

gzip is a compression utility used to compress the core files. Files created using the **gzip** command end with the file extension **.gz**. Compress the core files if they are larger than 50 MB.

Action

To compress the core files with gzip, use the following command from the shell:

```
root@host% gzip filename
```

Sample Output

```
root@host% gzip core-SCB.101010920452
```

Meaning

The contents of the core file are compressed into a single compressed file named **core-SCB.10101092045.gz**. The **gzip** command preserves the mode, ownership, and timestamps of files when compressing or decompressing them.

Log Software Version Information

Purpose

To log the Junos OS version information.

Action

To log the Junos OS version information, use the following Junos OS CLI operational mode command:

```
user@host> show version
```

Sample Output

```
user@host> show version
```

```
Hostname: host
Model: m10
JUNOS Base OS boot [5.0R5]
JUNOS Base OS Software Suite [5.0R5]
JUNOS Kernel Software Suite [5.0R5]
JUNOS Routing Software Suite [5.0R5]
JUNOS Packet Forwarding Engine Support [5.0R5]
JUNOS Crypto Software Suite [5.0R5]
JUNOS Online Documentation [5.0R5]
```



```

KERNEL 5.0R5 #0 built by builder on 2002-03-02 05:10:28 UTC
MGD release 5.0R5 built by builder on 2002-03-02 04:45:32 UTC
CLI release 5.0R5 built by builder on 2002-03-02 04:44:22 UTC
CHASSISD release 5.0R5 built by builder on 2002-03-02 04:43:37 UTC
DCD release 5.0R5 built by builder on 2002-03-02 04:42:47 UTC
RPD release 5.0R5 built by builder on 2002-03-02 04:46:17 UTC
SNMPD release 5.0R5 built by builder on 2002-03-02 04:52:26 UTC
MIB2D release 5.0R5 built by builder on 2002-03-02 04:45:37 UTC
APSD release 5.0R5 built by builder on 2002-03-02 04:43:31 UTC
VRRPD release 5.0R5 built by builder on 2002-03-02 04:52:34 UTC
ALARMD release 5.0R5 built by builder on 2002-03-02 04:43:24 UTC
PFED release 5.0R5 built by builder on 2002-03-02 04:46:06 UTC
CRAFTD release 5.0R5 built by builder on 2002-03-02 04:44:30 UTC
SAMPLED release 5.0R5 built by builder on 2002-03-02 04:52:20 UTC
ILMID release 5.0R5 built by builder on 2002-03-02 04:45:21 UTC
BPRELAYD release 5.0R5 built by builder on 2002-03-02 04:42:41 UTC
RMOPD release 5.0R5 built by builder on 2002-03-02 04:46:11 UTC
jkerneld release 5.0R5 built by builder on 2002-03-02 04:41:07 UTC
jroute-dd release 5.0R5 built by builder on 2002-03-02 04:41:21 UTC
jdocs-dd release 5.0R5 built by builder on 2002-03-02 04:39:11 UTC

```

Meaning

The sample output shows the hostname, router model, and the different Junos OS packages, processes, and documents.

Open a Case with JTAC

Problem

Description: To open a case with JTAC, call one of the Support phone numbers or create a case via [Case Manager](#).

Display the Crash Stack Traceback and Registration Information

Purpose

To display the crash stack traceback and registration information.

Action

To display the crash stack traceback and registration information, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:


```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Establish a vty session to the appropriate component. Use the **vtty** command followed by the executable name for the component; for example, **scb**, **ssb0**, **ssb1**, **fpc0**, or **fpc1**:

```
root@host% vty sfm0
```

NOTE: For the M40e and M160 routers, you can also create a **cty** session to the components if the components are not online.

4. Type the **show nvram** command to view the NVRAM information.
5. Type the **show syslog messages** command to view the system log messages.

Sample Output 1

```
user@host> start shell
```

```
% su
Password: ****
root@host% vty sfm0

SFM platform (266Mhz PPC 603e processor, 64Mb memory, 512Kb flash)
SFM3(host vty)# show nvram
System NVRAM :
  4080 available bytes, 4080 used, 0 free
Contents:

mpc106 machine check caused by error on the PCI Bus
mpc106 error detect register 1: 0x08, 2: 0x00
```



```

mpc106 error ack count = 0
mpc106 error address: 0x0a000000
mpc106 PCI bus error status register: 0x02
    mpc106 was the PCI master
    C/BE bits: I/O read [0b0010]
mpc106 error detection reg1: PCI cycle
mpc106 PCI status reg: parity error

```

```

System Exception: Vector/Code 0x00700, Signal 4
Event occurred at: Oct 26 13:32:40.952

```

```

Juniper Embedded Microkernel Version 4.2R1
Built by tlim on 2000-09-23 06:11:28 UTC
Copyright (C) 1998-2000, Juniper Networks, Inc.
All rights reserved.
Reason string: "Program Check"
Context: Thread (PFE Manager)

```

Registers:

```

R00: 0x06f5f81c R01: 0x06f5f9cc R02: 0x00003344 R03: 0x00000000
R04: 0x00008000 R05: 0x00000000 R06: 0x0010052c R07: 0x06f637e4
R08: 0x06f5f81c R09: 0x00169810 R10: 0x000000e8 R11: 0x00000001
R12: 0x00046cdf R13: 0xffffffff R14: 0xffffffff R15: 0xffffffff
R16: 0xffffffff R17: 0xffffffff R18: 0xffffffff R19: 0xffffffff
R20: 0xffffffff R21: 0xffffffff R22: 0xffffffff R23: 0xffffffff
R24: 0x00000003 R25: 0x00000000 R26: 0x00000001 R27: 0x0000fc78
R28: 0x00150000 R29: 0x0016c4b0 R30: 0x06f5eb7c R31: 0x97cb1d36
MSR: 0x0008b030 CTR: 0x000ac008 Link:0x06f5f81c SP: 0x06f5f9cc
CCR: 0x22200024 XER: 0x20000000 PC: 0x06f5f81c
DSISR: 0x00000000 DAR: 0xffffffff K_MSR: 0x00001030

```

Stack Traceback :

```

Frame 01: sp = 0x06f5f9cc, pc = 0x06f5f81c
Frame 02: sp = 0x06f5f9e4, pc = 0x000c7e28
Frame 03: sp = 0x06f5fa04, pc = 0x00026620

```

ROM NVRAM:

```
0 available bytes, 0 used, 0 free
```

```
SFM3(host vty)# show syslog messages
```

```

Oct 26 12:02:05 router tnp_sfm_2 PFEMAN: sent Resync request to Master
Oct 26 12:02:07 router tnp_sfm_3 CM(3): Slot 1: On-line
Oct 26 12:02:07 router tnp_sfm_3 CM(3): Slot 2: On-line

```



```

Oct 26 12:02:07 router tnp_sfm_3 CM(3): Slot 6: On-line
Oct 26 12:02:07 router tnp_sfm_3 PFEMAN: sent Resync request to Master
Oct 26 12:05:58 router tnp_sfm_3 mpc106 machine check caused by error on the
PCI Bu
s
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error detect register 1: 0x08,
2: 0x00
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error ack count = 0
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error address: 0x0a000000
Oct 26 12:05:58 router tnp_sfm_3 mpc106 PCI bus error status register: 0x02
Oct 26 12:05:58 router tnp_sfm_3 mpc106 was the PCI master
Oct 26 12:05:58 router tnp_sfm_3 C/BE bits: I/O read [0b0010]
Oct 26 12:05:58 router tnp_sfm_3 mpc106 error detection reg1: PCI cycle
Oct 26 12:05:58 router tnp_sfm_3 mpc106 PCI status reg: parity error
Oct 26 12:05:58 router tnp_sfm_3 ^B
Oct 26 12:05:58 router tnp_sfm_3 last message repeated 7 times
Oct 26 12:05:58 router tnp_sfm_3 Registers:
Oct 26 12:05:58 router tnp_sfm_3 R00: 0x06f5f81c R01: 0x06f5f9cc
R02: 0x00003344 R0
3: 0x00000000
Oct 26 12:05:58 router tnp_sfm_3 R04: 0x00008000 R05: 0x00000000
R06: 0x0010052c R0
7: 0x06f637e4
Oct 26 12:05:58 router tnp_sfm_3 R08: 0x06f5f81c R09: 0x00169810
R10: 0x000003b4 R1
1: 0x00000001
Oct 26 12:05:58 router tnp_sfm_3 R12: 0x00017b97 R13: 0xffffffff
R14: 0xffffffff R1
5: 0xffffffff
Oct 26 12:05:58 router tnp_sfm_3 R16: 0xffffffff R17: 0xffffffff
R18: 0xffffffff R1
9: 0xffffffff
Oct 26 12:05:58 router tnp_sfm_3 R20: 0xffffffff R21: 0xffffffff
R22: 0xffffffff R2
3: 0xffffffff
Oct 26 12:05:58 router tnp_sfm_3 R24: 0x00000003 R25: 0x00000000
R26: 0x00000001 R2
7: 0x0000fc78
Oct 26 12:05:58 router tnp_sfm_3 R28: 0x00150000 R29: 0x0016c4b0
R30: 0x06f5eb7c R3
1: 0x97c9c35e
Oct 26 12:05:58 router tnp_sfm_3 MSR: 0x0008b030 CTR: 0x000ac008
Link:0x06f5f81c SP
: 0x06f5f9cc

```



```
Oct 26 12:05:58 router tnp_sfm_3 CCR: 0x22200024 XER: 0x20000000
PC: 0x06f5f81c
Oct 26 12:05:58 router tnp_sfm_3 DSISR: 0x00000000 DAR: 0xffffffff
K_MSR: 0x0000103
0
```

Sample Output

The following sample output is another example of displaying the crash stack traceback and registration information:

```
root@host% vty fpc1
FPC160 platform (PPC 603e processor, 32Mb memory, 512Kb flash)

FPC1(host vty)# show nvram
System NVRAM :
  4080 available bytes, 4080 used, 0 free
  Contents:
00000000 R06: 0x0000005c R07: 0x850400d0
R08: 0x00000000 R09: 0x00000020 R10: 0x00000000 R11: 0x00000129
R12: 0x00000000 R13: 0x00000000 R14: 0x4005009a R15: 0x20000260
R16: 0xc8828784 R17: 0x84212800 R18: 0xc0004c61 R19: 0x80005900
R20: 0x80206000 R21: 0x84000304 R22: 0xd0410180 R23: 0x8c2005ac
R24: 0x00000003 R25: 0x00000000 R26: 0x00000001 R27: 0x0000fc48
R28: 0x001d0000 R29: 0x00000001 R30: 0x00136bb8 R31: 0x00000000
MSR: 0x0000b030 CTR: 0x001331e0 Link:0x000308c8 SP: 0x01baba34
CCR: 0x42200020 XER: 0x00000000 PC: 0x000308cc
DSISR: 0x00000000 DAR: 0xffffffff K_MSR: 0x00001030

Stack Traceback:
Frame 01: sp = 0x01baba34, pc = 0x000308c8
Frame 02: sp = 0x01babac4, pc = 0x0002647c
Frame 03: sp = 0x01babad4, pc = 0x00026590
Frame 04: sp = 0x01babadc, pc = 0x00106fcc
Frame 05: sp = 0x01babafc, pc = 0x00026620
ROM NVRAM:
  0 available bytes, 0 used, 0 free

FPC1(host vty)# show syslog messages
[0+00:00:00.780 LOG: Info] Version 4.0R5 by tlim on 2000-08-10 04:45:54 UTC
[0+00:00:00.780 LOG: Info] On-board NVRAM contains diagnostic information.
```



```
[0+00:00:03.175 LOG: Info] PFEMAN: Established connection to Master
[Jan 30 21:53:05.804 LOG: Info] SNTPD: Initial time of day set.
```

Meaning

Sample output 1 and 2 show the stack trace from the microkernel crash. Save the output from the **show nvram** and **show syslog** commands so that you can send them to JTAC when you open a case.

Clear the NVRAM Contents

Purpose

Currently the storage area for the logs on the NVRAM is limited to 4 KB. You need to delete old NVRAM logs to make room for new ones.

Action

To clear the content of the NVRAM after you have captured the necessary information, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```

2. Type **su** and the root password when prompted. You are now in the shell and the prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Establish a vty session to the appropriate component. Use the **vty** command followed by the abbreviation for the component, for example:

```
root@host% vty sfm0
SFM3(host vty)#
FPC1(host vty)#
```

4. Type the **clear nvram** command, for example:

```
SFM3(host vty)# clear nvram
FPC1(host vty)# clear nvram
```


Check Packet Forwarding Engine Microkernel Core Files

Purpose

If you observe disruption to the Packet Forwarding Engine microkernel, check the **/var/crash** directory for any core files created around the time of the crash.

Action

To check the **/var/crash** directory, use the following Junos OS CLI operational mode command:

```
user@host> file list detail /var/crash
```

Sample Output

```
user@host> file list detail /var/crash
```

```
var/crash:
total 456630
-rw-r--r--  1 root  wheel    6814720 Dec 18 08:03 core-FPC4.100111808032
-rw-r--r--  1 root  wheel    65613824 Dec 10 04:58 core-SCB.100111004570
-rw-r--r--  1 root  wheel    65613824 Dec 19 00:23 core-SCB.100111900221
-rw-r--r--  1 root  wheel    65545216 Feb  9 20:46 core-SCB.101010920452
```

Meaning

The sample output lists the contents of the **/var/crash/** directory. Check the date and timestamp for any core files created around the time of the crash. In the example above, four core files are listed.

List the Core Files Generated by the Crash

Purpose

To list the core files generated by the crash.

Action

To list the core files, follow these steps:

1. Exit from the CLI environment and create a UNIX-level shell by entering the **start shell** command:

```
user@host> start shell
```


2. Type **su** and the root password when prompted. You are now in the shell and the prompt is % instead of >, for example:

```
% su
Password: ****
```

3. Change the directory to **/var/crash** and type **ls -l**, for example:

```
root@host% cd /var/crash
root@host% ls -l
```

4. Look for any core files created around the time of the crash.

Sample Output

user@host> **start shell**

```
% su
Password: ****
root@host% cd /var/crash
root@host% ls -l
total 456630
-rw-r--r--  1 root  wheel    6814720 Dec 18 08:03 core-FPC4.100111808032
-rw-r--r--  1 root  wheel    65613824 Dec 10 04:58 core-SCB.100111004570
-rw-r--r--  1 root  wheel    65613824 Dec 19 00:23 core-SCB.100111900221
-rw-r--r--  1 root  wheel    65545216 Feb  9 20:46 core-SCB.101010920452
```

Meaning

The sample output shows the current core files for the different components on the router; for example, **core-FPC4.100111808032** and **core-SCB.100111004570**.

Compress the Core Files

Purpose

gzip is a compression utility used to compress the core files. Files created using the **gzip** command end with the file extension **.gz**. Compress the core files if they are larger than 50 MB.

Action

To compress the core files with **gzip**, use the following command from the shell:


```
root@host% gzip filename
```

Sample Output

```
root@host% gzip core-SCB.101010920452
```

Meaning

The contents of the core file are compressed into a single compressed file named **core-SCB.10101092045.gz**. The **gzip** command preserves the mode, ownership, and timestamps of files when compressing or decompressing them.

Configure a Primary Path

Action

To configure a primary path with an ERO list, bandwidth, and priority, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols mpls
```

2. Configure the primary ERO list:

```
[edit protocols mpls]
user@host# set path path-name address strict
```

For example:

```
[edit protocols mpls]
user@R1# set path via-r2 10.0.12.14 strict
user@R1# set path via-r2 10.0.24.2 strict
```

3. Configure the LSP:

```
[edit protocols mpls]
user@host# set label-switched-path lsp-path-name to destination;
```

For example:


```
[edit protocols mpls]
user@R1# set label-switched-path r1-to-r5 to 192.168.5.1;
```

4. Configure the primary path:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@host# set primary primary-name
```

For example:

```
[edit protocols mpls label-switched-path r1-to-r5]
user@R1# set primary via-r2
```

5. Configure the bandwidth:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@host# set primary primary-name bandwidth bandwidth
```

For example:

```
[edit protocols mpls label-switched-path r1-to-r5]
user@R1# set primary via-r2 bandwidth 35m
```

6. Configure the priority value:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@host# set primary primary-name priority reservation-priority setup-priority
```

For example:

```
[edit protocols mpls label-switched-path r1-to-r5]
user@R1# set primary via-r2 priority 6 6
```

7. Verify and commit the configuration:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@host# show
user@host# commit
```

Sample Output

The sample output below illustrates the configuration of the primary path on ingress router **R1**.


```

[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r5 {
  to 192.168.5.1;
  primary via-r2 { # Bandwidth and priority configured at the primary path
    bandwidth 35m; # level of the hierarchy
    priority 6 6; # Priority setup and hold values
  }
}
path via-r2 { # Primary ERO list
  10.0.12.14 strict;
  10.0.24.2 strict;
[...Output truncated...]

[edit protocols mpls]
user@R1# commit
commit complete

```

Meaning

The sample output shows a label-switched path (LSP) with bandwidth and priority applied to only one primary path. The same parameters specified one level up in the hierarchy, at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level, affect all paths.

The path, **via-r2**, specifies the complete strict path from the ingress to the egress routers through **10.0.12.14**, **10.0.24.2**, in that order. There cannot be any intermediate routers except the ones specified. However, there can be intermediate routers between **10.0.24.2** and the egress router because the egress router is not specifically listed in the path statement. To prevent intermediate routers before egress, configure the egress router as the last router, with a strict type.

Ensuring That Secondary Paths Establish When Resources Are Diminished

The Junos OS does not require that a primary and secondary path share the same parameters. You may decide to configure your primary paths with strict resource requirements, and configure your secondary paths with less strict requirements, allowing your secondary paths to establish more readily during periods of diminished resources.

Action

To ensure that secondary paths establish when resources are diminished, follow these steps:

1. In configuration mode, go to the following hierarchy level:


```
[edit protocols mpls]
user@host# edit label-switched-path lsp-path-name
```

For example:

```
[edit protocols mpls]
user@R1# edit label-switched-path r1-to-r4
```

2. Configure the bandwidth for the primary path, and do not configure any bandwidth for the secondary path:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@host# set primary primary-name bandwidth bandwidth
```

For example:

```
[edit protocols mpls label-switched-path r1-to-r4]
user@R1# set primary via-r2 bandwidth 35m
```

3. Verify and commit the configuration:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@host# show
user@host# commit
```

Sample Output

The sample output below illustrates a bandwidth configuration on ingress router R1.

```
[edit protocols mpls]
user@R1# show
label-switched-path r1-to-r4 {
  to 192.168.4.1;
  primary via-r2 {
    bandwidth 35m;
  }
  secondary via-r7 { # In this example, bandwidth is not configured for the
secondary path.
    standby;        # However you could configure a bandwidth value different
from
  }                  # that on the primary path.
```



```
}
[...Output truncated...]
```

Meaning

The sample output shows the primary path **via-r2** requires 35 Mbps of bandwidth, while secondary path **via-r7** has no constraints. The primary path is configured with strict resource requirements, while the secondary path is configured with no bandwidth requirements, allowing the secondary path to establish more readily during periods of diminished resources. One thing to keep in mind when configuring a secondary path without bandwidth requirements is that it can be subject to traffic loss due to congestion.

One-to-One Backup Overview

Fast reroute or one-to-one backup is a short-term solution to reduce packet loss associated with a particular LSP. One-to-one backup is appropriate under the following circumstances:

- Protection of a small number of LSPs relative to the total number of LSPs.
- Path selection criteria, such as bandwidth, priority, and link coloring for detour paths is critical.
- Control of individual LSPs is important.

In one-to-one backup, the ingress router adds the fast reroute object to the RSVP Path message requesting that downstream routers establish detours. Downstream routers generate Path messages and establish detours to avoid the downstream link or node. Detours are always calculated to avoid the immediate downstream link and node, providing against both link and node failure, as shown in [Figure 163 on page 2163](#).

Figure 163: One-to-One Backup Detours

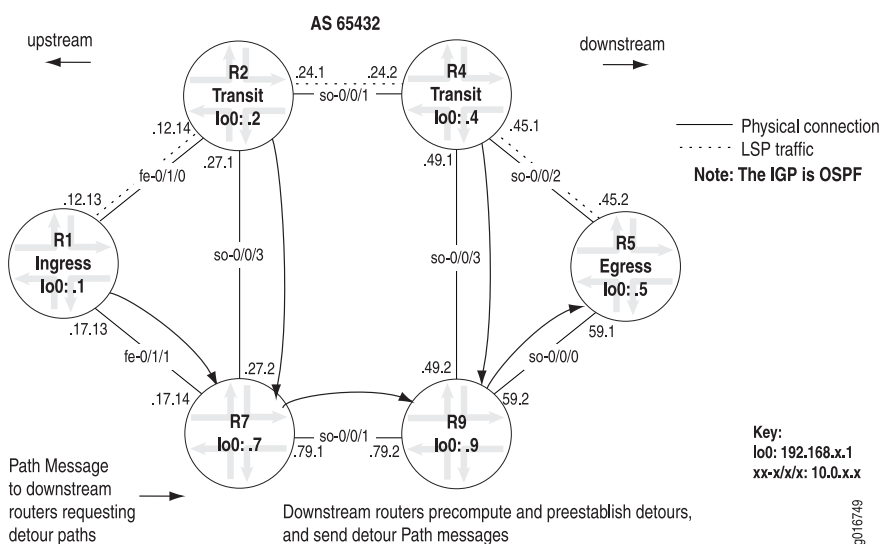


Figure 163 on page 2163 shows a network with one LSP configured from the ingress router **R1** to the egress router **R5**, transiting **R2** and **R4**. The following detours are established:

- **R1** creates a detour to **R5** via **R7** and **R9**
- **R2** creates a detour to **R5** via **R7** and **R9**
- **R4** creates a detour to **R5** via **R9**

Each detour is dedicated to a particular LSP traversing the router (one detour to one LSP). If the network topology has insufficient links and nodes, it may be impossible to establish a detour. Also, detour paths are not meant for long-term use because they may provide inadequate bandwidth and can result in congestion on the links. As soon as the ingress router calculates a new path avoiding the failure, traffic is redirected along the new path, detours are torn down, and new detours established.

Configure Link Protection

Purpose

Configuring link protection is a two-part process. The first part involves configuring link protection on the RSVP interface, and the second part sets link protection for any LSPs traversing the protected link that require use of the bypass path.

Action

To configure link protection, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@R1# edit protocols rsvp interface type-fpc/pic/port
```

For example:

```
[edit]
user@R1# edit protocols rsvp interface fe-0/1/0
```

2. Configure link protection for the interface:

```
[edit protocols rsvp interface type-fpc/pic/port]
user@R1# set link-protection
```

3. Verify the link protection configuration for the interface:

```
[edit protocols rsvp interface type-fpc/pic/port]
```



```
user@R1# show
```

4. Configure link protection for LSPs requiring use of the bypass path:

```
[edit protocols rsvp interface fe-0/1/0.0]
user@R1# top
[edit]
user@R1# edit protocols mpls label-switched-path lsp-path-name
```

For example:

```
[edit]
user@R1# edit protocols mpls label-switched-path lsp2-r1-to-r5
```

5. Configure link protection for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# set link-protection
```

6. Verify and the link protection configuration for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# show
user@R1# commit
```

Sample Output

The following sample output illustrates the configuration of the link protection on ingress router R1 in the network shown in [Figure 147 on page 1906](#):

```
[edit protocols rsvp]
user@R1# show
interface fe-0/1/0.0 {
    link-protection; #Protection for the RSVP interface
}

[edit protocols mpls label-switched-path lsp2-r1-to-r5]
user@R1# up

[edit protocols mpls]
user@R1# show
```



```

label-switched-path lsp2-r1-to-r5 { #Path level of the hierarchy
    #to 192.168.5.1;
    link-protection;

[edit protocols mpls]
user@R1# commit
commit complete

```

Meaning

The sample output shows link protection for a specific interface. After link protection is configured, a bypass path is signaled to avoid that link in case of a failure. Having a bypass path available does not in itself provide protection for LSPs that traverse the protected link. You must configure link protection on the ingress router for each LSP that will benefit from the bypass path.

Configuring and Verifying Link Protection

The following sections describe the steps you must take to configure and verify link protection (many-to-one backup):

1. [Configure Link Protection | 2166](#)
2. [Verify That Link Protection Is Up | 2168](#)

Configure Link Protection

Purpose

Configuring link protection is a two-part process. The first part involves configuring link protection on the RSVP interface, and the second part sets link protection for any LSPs traversing the protected link that require use of the bypass path.

Action

To configure link protection, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```

[edit]
user@R1# edit protocols rsvp interface type-fpc/pic/port

```

For example:


```
[edit]
user@R1# edit protocols rsvp interface fe-0/1/0
```

2. Configure link protection for the interface:

```
[edit protocols rsvp interface type-fpc/pic/port]
user@R1# set link-protection
```

3. Verify the link protection configuration for the interface:

```
[edit protocols rsvp interface type-fpc/pic/port]
user@R1# show
```

4. Configure link protection for LSPs requiring use of the bypass path:

```
[edit protocols rsvp interface fe-0/1/0.0]
user@R1# top
[edit]
user@R1# edit protocols mpls label-switched-path lsp-path-name
```

For example:

```
[edit]
user@R1# edit protocols mpls label-switched-path lsp2-r1-to-r5
```

5. Configure link protection for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# set link-protection
```

6. Verify and the link protection configuration for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# show
user@R1# commit
```

Sample Output

The following sample output illustrates the configuration of the link protection on ingress router R1 in the network shown in [Figure 147 on page 1906](#):


```

[edit protocols rsvp]
user@R1# show
interface fe-0/1/0.0 {
    link-protection; #Protection for the RSVP interface
}

[edit protocols mpls label-switched-path lsp2-r1-to-r5]
user@R1# up

[edit protocols mpls]
user@R1# show
label-switched-path lsp2-r1-to-r5 { #Path level of the hierarchy
    #to 192.168.5.1;
    link-protection;

[edit protocols mpls]
user@R1# commit
commit complete

```

Meaning

The sample output shows link protection for a specific interface. After link protection is configured, a bypass path is signaled to avoid that link in case of a failure. Having a bypass path available does not in itself provide protection for LSPs that traverse the protected link. You must configure link protection on the ingress router for each LSP that will benefit from the bypass path.

Verify That Link Protection Is Up

Purpose

When you verify link protection, you must check that the bypass LSP is up. You can also check the number of LSPs protected by the bypass. In the network shown in [Figure 147 on page 1906](#), a bypass path should be up to protect the link between **R1** and **R2**, or next-hop **10.0.12.14**, and the two LSPs traversing the link, **lsp2-r1-to-r5** and **lsp1-r6-to-r0**.

Action

To verify link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router:

```

user@host> show mpls lsp extensive
user@host> show rsvp session detail
user@host> show rsvp interface

```


Sample Output

```

user@R1> show mpls lsp extensive | no-more
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: lsp2-r1-to-r5
  ActivePath:   via-r2 (primary)
Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary   via-r2           State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
      10.0.12.14(Label=101264) 10.0.24.2(Label=100736) 10.0.45.2(Label=3)
    6 Jun 16 14:06:33   Link-protection Up
    5 Jun 16 14:05:39   Selected as active path
    4 Jun 16 14:05:39   Record Route:  10.0.12.14(Label=101264)
10.0.24.2(Label=100736) 10.0.45.2(Label=3)
    3 Jun 16 14:05:39   Up
    2 Jun 16 14:05:39   Originate Call
    1 Jun 16 14:05:39   CSPF: computation result accepted
    Created: Fri Jun 16 14:05:38 2006
Total 1 displayed, Up 1, Down 0

[...Output truncated...]

Transit LSP: 2 sessions

192.168.0.1
  From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
    LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101296
  Resv style: 1 SE, Label in: 100192, Label out: 101296
  Time left: 116, Since: Mon Jun 19 10:26:32 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 58739 protocol 0
Link protection desired
  Type: Link protected LSP, using Bypass->10.0.12.14
    1 Jun 19 10:26:32 Link protection up, using Bypass->10.0.12.14

```



```

PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 579 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.12.14 (fe-0/1/0.0) 474 pkts
RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 501 pkts
Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
[...Output truncated...]

```

Meaning

The sample output from ingress router **R1** shows that **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have link protection up, and both LSPs are using the bypass path, **10.0.12.14**. However, the **show mpls lsp** command does not list the bypass path. For information about the bypass path, use the **show rsvp session** command.

Sample Output

```

user@R1> show rsvp session detail
Ingress RSVP: 2 sessions
192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101456
  Resv style: 1 SE, Label in: -, Label out: 101456
  Time left: -, Since: Fri May 26 18:38:09 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 18709 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 51939 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 55095 pkts
  Explct route: 10.0.17.14 10.0.27.1
  Record route: <self> 10.0.17.14 10.0.27.1

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp2-r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101264

```


Resv style: 1 SE, Label in: -, Label out: 101264
 Time left: -, Since: Fri Jun 16 14:05:39 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 18724 protocol 0

Link protection desired

Type: Link protected LSP

PATH rcvfrom: localclient
 Adspec: sent MTU 1500
 Path MTU: received 1500
 PATH sentto: 10.0.12.14 (fe-0/1/0.0) 8477 pkts
 RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 8992 pkts
 Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
 Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
 Total 2 displayed, Up 2, Down 0

Egress RSVP: 1 sessions

192.168.1.1

From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0
 LSPname: r5-to-r1, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: -
 Resv style: 1 FF, Label in: 3, Label out: -
 Time left: 159, Since: Mon May 22 22:08:16 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 64449 protocol 0
 PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 63145 pkts
 Adspec: received MTU 1500
 PATH sentto: localclient
 RESV rcvfrom: localclient
 Record route: 10.0.59.1 10.0.79.2 10.0.17.14 <self>

Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions

192.168.0.1

From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
LSPname: lsp1-r6-to-r0, LSPpath: Primary
 Suggested label received: -, Suggested label sent: -
 Recovery label received: -, Recovery label sent: 101296
 Resv style: 1 SE, Label in: 100192, Label out: 101296
 Time left: 129, Since: Mon Jun 19 10:26:32 2006
 Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
 Port number: sender 1 receiver 58739 protocol 0

Link protection desired

Type: Link protected LSP


```

PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 3128 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.0.12.14 (fe-0/1/0.0) 2533 pkts
RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 2685 pkts
Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-r6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100128, Label out: 3
  Time left: 143, Since: Thu May 25 12:30:26 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 4111 protocol 0
  PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 57716 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 54524 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 50534 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.59.1 10.0.79.2 10.0.17.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning

The sample output from ingress router **R1** shows the ingress, egress, and transit LSPs for **R1**. Some information is similar to that found in the **show mpls lsp** command. However, because link protection is an RSVP feature, information about bypass paths is provided. The bypass path appears as a separate RSVP ingress session for the protected interface, as indicated by the **Type** field.

The bypass path name is automatically generated. By default, the name appears as **Bypass > interface-address**, where the interface address is the next downstream router's interface (**10.0.12.14**). The explicit route **10.0.17.14 10.0.27.1** for the session shows **R7** as the transit node and **R2** as the egress node.

Within the ingress RSVP section of the output, the LSP originating at **R1** (**lsp2-r1-to-r5**) is shown requesting link protection. Since a bypass path is in place to protect the downstream link, **lsp2-r1-to-r5** is associated with the bypass, as indicated by the **Link protected LSP** field.

The egress section of the output shows the return LSP **r5-to-r1**, which is not protected.

The transit section of the output shows link protection requested by **lsp1-r6-to-r0**. Since a bypass path is in place to protect the downstream link, **lsp1-r6-to-r0** is associated with the bypass, as indicated by the **Link protected LSP** field. Also included in the transit section of the output is the return LSP **r0-to-r6**, which is not protected.

Sample Output

```
user@R1> show rsvp interface
RSVP interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
fe-0/1/0.0	Up	2	100%	100Mbps	100Mbps	0bps	35Mbps
fe-0/1/1.0	Up	1	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/2.0	Up	0	100%	100Mbps	100Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning

The sample output from ingress router **R1** shows the number of LSPs going through the interfaces configured on **R1**. The **Active resv** field shows the number of LSPs for each interface. For example, interface **fe-0/1/0.0** between **R1** and **R2** has two active reservations, **lsp1-r6-to-r0** and **lsp2-r1-to-r5**; interface **fe-0/1/1.0** between **R1** and **R7** has one, the bypass (**10.0.12.14**); interface **fe-0/1/2.0** between **R6** and **R1** has no LSP reservations; and interface **so-0/0/3.0** between **R6** and **R1** has one LSP reservation, **lsp1-r6-to-r0**.

Configure Node-Link Protection

Configuring node-link protection is a two-part process. The first part involves configuring node-link protection for any LSPs traversing the protected node that require use of the bypass path, and the second part sets link protection on the outgoing RSVP interface on routers in the LSP.

Action

To configure node-link protection, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@R1# edit protocols mpls label-switched-path lsp-path-name
```

For example:

```
[edit]
user@R1# edit protocols mpls label-switched-path lsp2-r1-to-r5
```

2. Configure node-link protection for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
```



```
user@R1# set node-link-protection
```

3. Verify the node-link protection configuration for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# show
```

4. Configure link protection for the interface:

```
[edit protocols]
user@R1# edit protocols rsvp interface interface-name
```

For example:

```
[edit protocols]
user@R1# edit protocols rsvp interface fe-0/1/0
```

5. Configure link protection:

```
[edit protocols rsvp interface interface-name]
user@R1# set link-protection
```

6. Verify the link protection configuration for the interface, and commit both configurations:

```
[edit protocols rsvp interface interface-name]
user@R1# show
user@R1# commit
```

7. Repeat Step 1 through Step 3 on any other ingress routers that have LSPs requiring use of the bypass path.
8. Repeat Step 4 and Step 5 on routers with outgoing interfaces in the LSP.

Sample Output

The following sample output shows the configuration of node-link protection on ingress router R1 in the network shown in [Figure 147 on page 1906](#):

```
[edit protocols mpls label-switched-path lsp2-r1-to-r5]
user@R1# up

[edit protocols mpls]
```



```

user@R1# show
label-switched-path lsp2-r1-to-r5 { #Label-switched-path level of the hierarchy

    to 192.168.5.1;
    node-link-protection; #LSP node-link protection

[edit protocols rsvp]
user@R1# show
interface fe-0/1/0.0 {
    link-protection; #Link protection for the RSVP interface
}

[edit protocols rsvp]
user@R1# commit
commit complete

```

Meaning

The sample output shows the configuration of node-link protection for an LSP. After node-link protection is configured, bypass paths are signaled to avoid the protected link or node in case of failure. Having bypass paths available does not in itself provide protection for LSPs that traverse the protected node. You must include the **node-link-protection** statement on the ingress router for each LSP that will benefit from the bypass path.

Configuring and Verifying Node-Link Protection

The following section describes the steps you must take to configure and verify many-to-one backup.

1. [Configure Node-Link Protection | 2175](#)
2. [Verify That Node-Link Protection Is Up | 2177](#)

Configure Node-Link Protection

Configuring node-link protection is a two-part process. The first part involves configuring node-link protection for any LSPs traversing the protected node that require use of the bypass path, and the second part sets link protection on the outgoing RSVP interface on routers in the LSP.

Action

To configure node-link protection, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@R1# edit protocols mpls label-switched-path lsp-path-name
```

For example:

```
[edit]
user@R1# edit protocols mpls label-switched-path lsp2-r1-to-r5
```

2. Configure node-link protection for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# set node-link-protection
```

3. Verify the node-link protection configuration for the LSP:

```
[edit protocols mpls label-switched-path lsp-path-name]
user@R1# show
```

4. Configure link protection for the interface:

```
[edit protocols]
user@R1# edit protocols rsvp interface interface-name
```

For example:

```
[edit protocols]
user@R1# edit protocols rsvp interface fe-0/1/0
```

5. Configure link protection:

```
[edit protocols rsvp interface interface-name]
user@R1# set link-protection
```

6. Verify the link protection configuration for the interface, and commit both configurations:

```
[edit protocols rsvp interface interface-name]
user@R1# show
user@R1# commit
```


7. Repeat Step 1 through Step 3 on any other ingress routers that have LSPs requiring use of the bypass path.
8. Repeat Step 4 and Step 5 on routers with outgoing interfaces in the LSP.

Sample Output

The following sample output shows the configuration of node-link protection on ingress router R1 in the network shown in [Figure 147 on page 1906](#):

```
[edit protocols mpls label-switched-path lsp2-r1-to-r5]
user@R1# up

[edit protocols mpls]
user@R1# show
label-switched-path lsp2-r1-to-r5 { #Label-switched-path level of the hierarchy

    to 192.168.5.1;
    node-link-protection; #LSP node-link protection

[edit protocols rsvp]
user@R1# show
interface fe-0/1/0.0 {
    link-protection; #Link protection for the RSVP interface
}

[edit protocols rsvp]
user@R1# commit
commit complete
```

Meaning

The sample output shows the configuration of node-link protection for an LSP. After node-link protection is configured, bypass paths are signaled to avoid the protected link or node in case of failure. Having bypass paths available does not in itself provide protection for LSPs that traverse the protected node. You must include the **node-link-protection** statement on the ingress router for each LSP that will benefit from the bypass path.

Verify That Node-Link Protection Is Up

Purpose

After you configure node-link protection, you must check that bypass paths are up. You can also check the number of LSPs protected by the bypass paths. In the network shown in [Figure 7 on page 144](#), two

bypass paths should be up: one next-hop bypass path protecting the link between **R1** and **R2** (or next-hop **10.0.12.14**), and a next-next-hop bypass path avoiding **R2**.

Action

To verify node-link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router. You can also issue the commands on transit routers and other routers used in the bypass path for slightly different information.

```
show mpls lsp
show mpls lsp extensive
show rsvp interface
show rsvp interface extensive
show rsvp session detail
```

Sample Output

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To                From                State Rt ActivePath      P      LSPname
192.168.5.1       192.168.1.1      Up    0  via-r2          *      lsp2-r1-to-r5
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions
To                From                State Rt Style Labelin Labelout LSPname
192.168.1.1       192.168.5.1      Up    0  1 FF           3      - r5-to-r1
Total 1 displayed, Up 1 , Down 0

Transit LSP: 2 sessions
To                From                State Rt Style Labelin Labelout LSPname
192.168.0.1       192.168.6.1      Up    0  1 FF 100464 101952  lsp1-r6-to-r0
192.168.6.1       192.168.0.1      Up    0  1 FF 100448      3  r0-to-t6
Total 2 displayed, Up 2, Down 0
```

Meaning

Sample output from **R1** for the **show mpls lsp** command shows a brief description of the state of configured and active LSPs for which **R1** is the ingress, transit, and egress router. All LSPs are up. **R1** is the ingress router for **lsp2-r1-to-r5**, and the egress router for return LSP **r5-to-r1**. Two LSPs transit **R1**, **lsp1-r6-to-r0** and the return LSP **r0-to-t6**. For more detailed information about the LSP, include the **extensive** option when you issue the **show mpls lsp** command.

Sample Output

```

user@R1> show mpls lsp extensive

Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up , ActiveRoute: 0, LSPname: lsp2-r1-to-r5
  ActivePath: via-r2 (primary)
    Node/Link protection desired
    LoadBalance: Random
    Encoding type: Packet, Switching type: Packet, GPID: IPv4
    *Primary via-r2 State: Up
      SmartOptimizeTimer: 180
      Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      10.0.12.14(Label=101872) 10.0.24.2(Label=101360) 10.0.45.2(Label=3)
    11 Jul 11 14:30:58 Link-protection Up
    10 Jul 11 14:28:28 Selected as active path
    [...Output truncated...]
    Created: Tue Jul 11 14:22:58 2006
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.1.1
  From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0
    LSPname: r5-to-r1, LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: -
    Resv style: 1 FF, Label in: 3, Label out: -
    Time left: 146, Since: Tue Jul 11 14:28:36 2006
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 1 receiver 29228 protocol 0
    PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 362 pkts
    Adspec: received MTU 1500
    PATH sentto: localclient
    RESV rcvfrom: localclient
    Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 2 sessions

```



```

192.168.0.1
  From: 192.168.6.1,  LSPstate: Up, ActiveRoute: 0
    LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101952
  Resv style: 1 SE, Label in: 100464, Label out: 101952
  Time left: 157, Since: Tue Jul 11 14:31:38 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 11131 protocol 0
    Node/Link protection desired
    Type: Node/Link protected LSP, using Bypass->10.0.12.14->10.0.24.2
      1 Jul 11 14:31:38 Node protection up, using Bypass->10.0.12.14->10.0.24.2
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 509 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 356 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1,  LSPstate: Up, ActiveRoute: 0
    LSPname: r0-to-t6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left: 147, Since: Tue Jul 11 14:31:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 23481 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 350 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 323 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning

Sample output from **R1** for the **show mpls lsp extensive** command shows detailed information about all LSPs for which **R1** is the ingress, egress, or transit router, including all past state history and the reason why an LSP failed. All LSPs are up. The main two LSPs **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have node-link protection as indicated by the **Node/Link protection desired** field in the ingress and transit sections of the output. In the ingress section of the output, the **Link-protection Up** field shows that **lsp2-r1-to-r5** has

link protection up. In the transit section of the output, the **Type: Node/Link protected LSP** field shows that **lsp1-r6-to-r0** has node-link protection up, and in case of failure will use the bypass LSP **Bypass->10.0.12.14->10.0.24.2**.

Sample Output

```
user@R1> show rsvp interface
RSVP interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
fe-0/1/0.0	Up	2	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/1.0	Up	1	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/2.0	Up	0	100%	100Mbps	100Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning

Sample output from **R1** for the **show rsvp interface** command shows four interfaces enabled with RSVP (**Up**). Interface **fe-0/1/0.0** has two active RSVP reservations (**Active resv**) that might indicate sessions for the two main LSPs, **lsp1-r6-to-r0** and **lsp2-r1-to-r5**. Interface **fe-0/1/0.0** is the connecting interface between R1 and R2, and both LSPs are configured with a strict path through **fe-0/1/0.0**. For more detailed information about what is happening on interface **fe-0/1/0.0**, issue the **show rsvp interface extensive** command.

Sample Output

```
user@R1> show rsvp interface extensive
RSVP interface: 3 active
fe-0/1/0.0 Index 67, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 10.0.12.13
ActiveResv 2, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 100Mbps
  MaxAvailableBW 100Mbps = (bc0*subscription)
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Protection: On, Bypass: 2, LSP: 2, Protected LSP: 2, Unprotected LSP: 0
  2 Jul 14 14:49:40 New bypass Bypass->10.0.12.14
  1 Jul 14 14:49:34 New bypass Bypass->10.0.12.14->10.0.24.2
Bypass: Bypass->10.0.12.14, State: Up, Type: LP, LSP: 0, Backup: 0
```



```

3 Jul 14 14:49:42 Record Route: 10.0.17.14 10.0.27.1
2 Jul 14 14:49:42 Up
1 Jul 14 14:49:42 CSPF: computation result accepted
Bypass: Bypass->10.0.12.14->10.0.24.2, State: Up, Type: NP, LSP: 2, Backup:0
4 Jul 14 14:50:04 Record Route: 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
3 Jul 14 14:50:04 Up
2 Jul 14 14:50:04 CSPF: computation result accepted
1 Jul 14 14:49:34 CSPF failed: no route toward 10.0.24.2
[...Output truncated...]

```

Meaning

Sample output from **R1** for the **show rsvp interface extensive** command shows more detailed information about the activity on all RSVP interfaces (3). However, only output for **fe-0/1/0.0** is shown. Protection is enabled (**Protection: On**), with two bypass paths (**Bypass: 2**) protecting two LSPs (**Protected LSP: 2**). All LSPs are protected, as indicated by the **Unprotected LSP: 0** field. The first bypass **Bypass->10.0.12.14** is a link protection bypass path (**Type: LP**), protecting the link between **R1** and **R2 fe-0/1/0.0**. The second bypass path **10.0.12.14->10.0.24.2** is a node-link protected LSP, avoiding **R2** in case of node failure.

Sample Output

```

user@R1> show rsvp session detail
Ingress RSVP: 2 sessions

192.168.4.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14->10.0.24.2
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102000
  Resv style: 1 SE, Label in: -, Label out: 102000
  Time left: -, Since: Tue Jul 11 14:30:53 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 60120 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 336 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 310 pkts
  Explot route: 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1

```


Record route: <self> 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1

192.168.5.1

From: 192.168.1.1, **LSPstate: Up**, ActiveRoute: 0

LSPname: lsp2-r1-to-r5, LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: 101872

Resv style: 1 SE, Label in: -, Label out: 101872

Time left: -, Since: Tue Jul 11 14:28:28 2006

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 2 receiver 60118 protocol 0

Node/Link protection desired

Type: Node/Link protected LSP

PATH rcvfrom: localclient

Adspec: sent MTU 1500

Path MTU: received 1500

PATH sentto: 10.0.12.14 (fe-0/1/0.0) 344 pkts

RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 349 pkts

Explt route: 10.0.12.14 10.0.24.2 10.0.45.2

Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2

Total 2 displayed, Up 2, Down 0

Egress RSVP: 1 sessions

192.168.1.1

From: 192.168.5.1, **LSPstate: Up**, ActiveRoute: 0

LSPname: r5-to-r1, LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: -

Resv style: 1 FF, Label in: 3, Label out: -

Time left: 147, Since: Tue Jul 11 14:28:36 2006

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 1 receiver 29228 protocol 0

PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 348 pkts

Adspec: received MTU 1500

PATH sentto: localclient

RESV rcvfrom: localclient

Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>

Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions

192.168.0.1


```

From: 192.168.6.1,  LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101952
  Resv style: 1 SE, Label in: 100464, Label out: 101952
  Time left: 134, Since: Tue Jul 11 14:31:38 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 11131 protocol 0
  Node/Link protection desired
  Type: Node/Link protected LSP
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 488 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 339 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 343 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1,  LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-t6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left: 158, Since: Tue Jul 11 14:31:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 23481 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 344 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 337 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 310 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning

Sample output from **R1** shows detailed information about the RSVP sessions active on **R1**. All sessions are up, with two ingress sessions, one egress session, and two transit sessions.

Within the ingress section, the first session is a bypass path, as indicated by the **Type: Bypass LSP** field; and the second session is a protected LSP (**lsp2-r1-to-r5**) originating on **R1**, as indicated by the **Type: Node/Link protected LSP** field.

Conclusion

Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection and node-link protection are facility-based methods used to reduce the amount of time needed to reroute LSP traffic. These protection methods are often compared to fast reroute—the other Junos OS LSP protection method.

While fast reroute protects LSPs on a one-to-one basis, link protection and node-link protection protect multiple LSPs by using a single, logical bypass LSP. Link protection provides robust backup support for a link, node-link protection bypasses a node or a link, and both types of protection are designed to interoperate with other vendor equipment. Such functionality makes link protection and node-link protection excellent choices for scalability, redundancy, and performance in MPLS-enabled networks.

Related Information

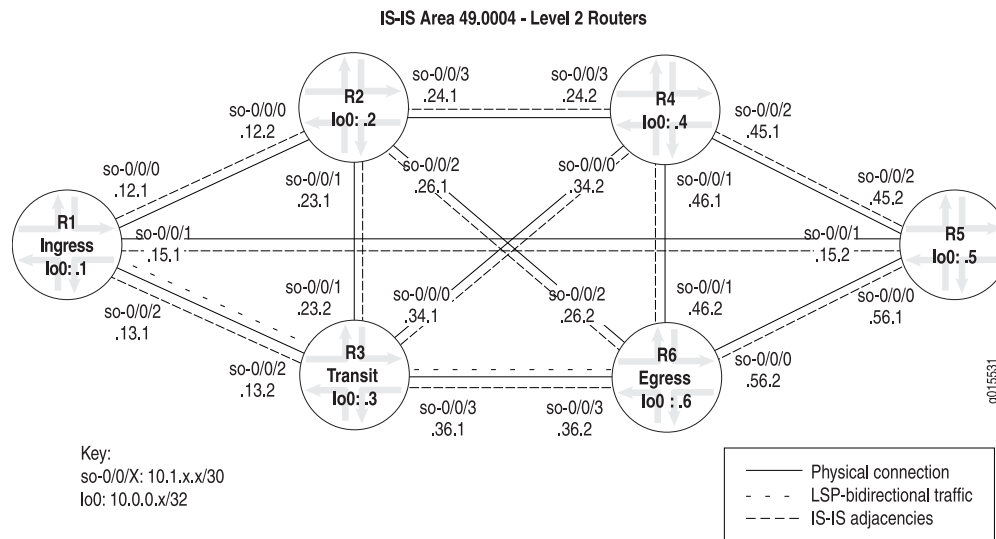
For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- *Junos User Guide*
- *Junos MPLS Applications Configuration Guide*
- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering*. White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

Configure IS-IS as the IGP

Before you can run MPLS on your network, you should have an IGP running on all specified routers and interfaces. The IGP can be either IS-IS or OSPF. For the steps to configure OSPF, see [“Configure OSPF as the IGP” on page 2196](#).

Figure 164: IS-IS Network Topology



The IS-IS IGP in the MPLS network in [Figure 164 on page 2186](#) consists of the following:

- All routers are configured for Level 2, therefore default CSPF LSPs can occur.
- All routers are in IS-IS area 49.0004. However, the routers in this network could be in any area because Level 2 adjacencies occur between all directly connected Level 2 routers regardless of which area they are in.
- Level 2 adjacencies between all directly connected Level 2 routers as follows:
 - R1 is adjacent to R2, R3, and R5
 - R2 is adjacent to R1, R3, R4, and R6
 - R3 is adjacent to R1, R2, R4, and R6
 - R4 is adjacent to R2, R3, R5, and R6
 - R5 is adjacent to R1, R4, and R6
 - R6 is adjacent to R2, R3, R4, and R5

When you configure IS-IS as the IGP, you must enable IS-IS on the router, configure International Organization for Standardization (ISO) addressing, and enable IS-IS on all router interfaces.

You can enable IS-IS throughout the rest of the network by repeating Step 1, “[Enable IS-IS on Routers in Your Network](#)” on page 2187 through Step 3, “[Enable IS-IS on Router Interfaces](#)” on page 2191 as appropriate on other routers until all routers and interfaces establish IS-IS adjacencies.

To configure IS-IS and establish IS-IS adjacencies, follow these steps:

1. [Enable IS-IS on Routers in Your Network](#) | 2187
2. [Configure ISO Addressing](#) | 2190

3. [Enable IS-IS on Router Interfaces | 2191](#)
4. [Verify That IS-IS Adjacencies Are Established | 2193](#)

Enable IS-IS on Routers in Your Network

Action

To enable IS-IS on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols isis
```

2. Disable Level 1 if appropriate for your network:

```
[edit protocols isis]
user@host# set level 1 disable
```

3. Configure the interface:

```
[edit protocols isis]
user@host# edit interface type-fpc/pic/port level level-number metric metric
```

4. Disable the management interface if you have included the **interface all** statement, as shown in [Sample Output 2 on page 2189](#):

```
[edit protocols isis]
user@host# set interface fxp0.0 disable
```

5. Include the loopback interface (**lo0**) if you have listed all interfaces separately, as shown in [Sample Output 1 on page 2188](#):

```
[edit protocols isis]
user@host# set interface lo0.0
```

6. Set the loopback interface (**lo0**) to passive:

```
[edit protocols isis]
user@R1# set interface lo0 passive
```

7. Verify and commit the configuration:


```
user@host# show
```

```
user@host# commit
```

Sample Output 1

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols isis

[edit protocols isis]
user@R1# set level 1 disable

[edit protocols isis]
user@host# edit interface all level 2 metric 10

[edit protocols isis]
user@host# set interface lo0.0

[edit protocols isis]
user@host# set interface lo0 passive

[edit protocols isis]
user@R1# show
level 1 disable;
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface lo0.0;
    passive;
}

[edit protocols isis]
user@R1# commit
commit complete
```


Sample Output 2

```
[edit protocols isis]
user@R6# show
level 1 disable;
interface all {
    level 2 metric 15;
}
interface fxp0.0 {
    disable;
}
interface lo0.0 {
    passive;
}
```

Meaning

Sample Output 1 shows that IS-IS Level 1 is disabled, making this a Level 2 router. When all routers in the network are running at one IS-IS level (Level 2), default CSPF LSPs can occur.

Because **R1** in Sample Output 1 has all IS-IS enabled interfaces listed, including the loopback interface (**lo0**), you do not need to include the **disable** statement for the management interface (**fxp0**). All interfaces have unit number **0**, the default if a unit number is not specified. When you configure an interface at the **[edit protocols isis]** hierarchy level, and you do not include the logical unit, the default **0** is appended to the interface name, for example, **so-0/0/1.0**.

Sample Output 2 does not list the interfaces configured with IS-IS; instead, all interfaces are configured, including the loopback interface (**lo0**) and the management interface (**fxp0**). Therefore, you do not need to include a separate statement for the loopback (**lo0**) interface. However, in this instance, it is best practice to disable the management interface (**fxp0**) so that IS-IS packets are not sent over it. If you do not disable the management interface (**fxp0**) when you include the **interface-all** statement, the IS-IS protocol can form adjacencies over the management backbone, but traffic does not flow because transit traffic does not go out of the management interface.

Sample Output 2 also shows that all interfaces on **R6** are configured with a metric of **15**. A metric is not required to configure IS-IS on your interfaces. The default metric value is **10** (with the exception of the loopback [**lo0**] interface, which has a default metric of **0**). A metric is included to demonstrate that you can configure a metric for IS-IS if the default (**10**) is not appropriate for your network.

Both sample outputs show the **passive** statement included in the configuration of the loopback (**lo0**) interface. Including the **passive** statement is considered best practice and ensures the following:

- Protocols are not run over the loopback (**lo0**) interface
- When the router ID (RID) is configured manually, ensures that the loopback (**lo0**) interface is advertised to other networks.

NOTE: It is considered best practice to configure the RID manually to avoid duplicate RID problems.

Configure ISO Addressing

Purpose

For a router to support IS-IS, you must configure an ISO network entity title (NET) address on one of the router's interfaces, preferably the loopback interface (**lo0**).

Action

To configure ISO addressing, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Include a NET address for the loopback interface:

```
[edit interfaces]
user@host# set lo0 unit number family iso address address
```

3. Verify and commit the configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set lo0 unit 0 family iso address 49.0004.1000.0000.0001.00
```



```

[edit interfaces]
userR1# show
[...Output truncated...]
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.1/32;
        }
        family iso {
            address 49.0004.1000.0000.0001.00;
        }
    }
}

[edit interfaces]
user@R1# commit
commit complete

```

Meaning

The sample output shows that the loopback (**lo0**) interface is configured with the NET address **49.0004.1000.0000.0001.00**. The loopback interface (**lo0**) becomes a point of connection from the router to the IS-IS network. Every router in an IS-IS network must have at least one ISO NET address that identifies a point of connection to the IS-IS network. The NET address is generally configured on the loopback (**lo0**) interface. Routers that participate in multiple areas can have multiple NET addresses.

All the routers in the network share a Level 2 database containing identical information. A common Level 2 database occurs in this case because all adjacencies are Level 2, and all routers are within the same IS-IS area (**49.0004**). Level 2 LSP flooding reaches all routers in the network due to the presence of a single level. For more information on determining the NET address, see the *Junos Routing Protocols Configuration Guide*.

Enable IS-IS on Router Interfaces

Purpose

Enable reception and transmission of ISO protocol data units (PDUs) on each router interface in the network with the **family** statement, which identifies which protocol packets are accepted into the interfaces. For example, valid IS-IS packets are dropped if the interface is not configured with the **family iso** statement.

Action

To configure support for IS-IS on router interfaces in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:


```
[edit]
user@host# edit interfaces
```

2. Configure IS-IS:

```
[edit interfaces]
user@host# set type-fpc/pic/port unit number family iso
```

3. Verify and commit the configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set so-0/0/2 unit 0 family iso

[edit interfaces]
user@R1# show
[...Output truncated...]
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.13.1/30;
        }
        family iso;
    }
}

[edit interfaces]
user@R1# commit
commit complete
```


Meaning

The sample output shows that the interface **so-0/0/2** is configured with IS-IS.

Verify That IS-IS Adjacencies Are Established**Purpose**

After configuring IS-IS, you must verify that neighboring routers have formed adjacencies with each other.

Action

To verify IS-IS adjacencies, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show isis adjacency
```

Sample Output

```
user@R1> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R2	2 Up	25	
so-0/0/1.0	R5	2 Up	23	
so-0/0/2.0	R3	2 Up	20	

```
user@R3> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R4	2 Up	25	
so-0/0/1.0	R2	2 Up	25	
so-0/0/2.0	R1	2 Up	26	
so-0/0/3.0	R6	2 Up	25	

```
user@R6> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R5	2 Up	19	
so-0/0/1.0	R4	2 Up	22	
so-0/0/2.0	R2	2 Up	22	
so-0/0/3.0	R3	2 Up	19	

Sample Output

```

user@R1> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R2          2 Up         25
so-0/0/1.0         R5          2 Up         23
so-0/0/2.0         R3          2 Up         20

user@R3> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R4          2 Up         25
so-0/0/1.0         R2          2 Up         25
so-0/0/2.0         R1          2 Up         26
so-0/0/3.0         R6          2 Up         25

user@R6> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R5          2 Up         19
so-0/0/1.0         R4          2 Up         22
so-0/0/2.0         R2          2 Up         22
so-0/0/3.0         R3          2 Up         19

```

Meaning

The sample output from the ingress, transit, and egress routers shows that all routers in the network have formed IS-IS adjacencies.

Verify That IS-IS Adjacencies Are Established

Purpose

After configuring IS-IS, you must verify that neighboring routers have formed adjacencies with each other.

Action

To verify IS-IS adjacencies, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show isis adjacency
```


Sample Output

```
user@R1> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R2	2 Up	25	
so-0/0/1.0	R5	2 Up	23	
so-0/0/2.0	R3	2 Up	20	

```
user@R3> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R4	2 Up	25	
so-0/0/1.0	R2	2 Up	25	
so-0/0/2.0	R1	2 Up	26	
so-0/0/3.0	R6	2 Up	25	

```
user@R6> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R5	2 Up	19	
so-0/0/1.0	R4	2 Up	22	
so-0/0/2.0	R2	2 Up	22	
so-0/0/3.0	R3	2 Up	19	

Sample Output

```
user@R1> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R2	2 Up	25	
so-0/0/1.0	R5	2 Up	23	
so-0/0/2.0	R3	2 Up	20	

```
user@R3> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R4	2 Up	25	
so-0/0/1.0	R2	2 Up	25	
so-0/0/2.0	R1	2 Up	26	
so-0/0/3.0	R6	2 Up	25	

```
user@R6> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
-----------	--------	---------	-------------	------

so-0/0/0.0	R5	2 Up	19
so-0/0/1.0	R4	2 Up	22
so-0/0/2.0	R2	2 Up	22
so-0/0/3.0	R3	2 Up	19

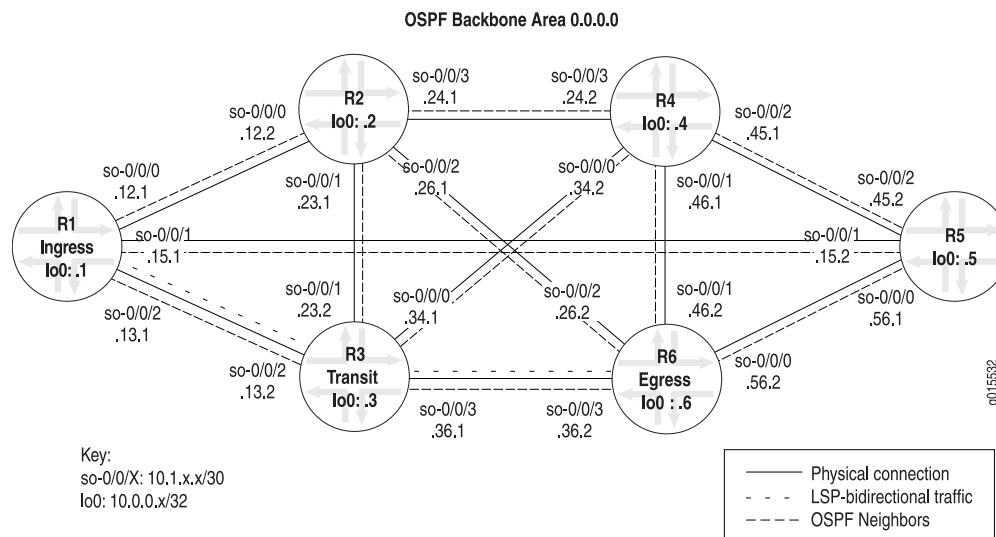
Meaning

The sample output from the ingress, transit, and egress routers shows that all routers in the network have formed IS-IS adjacencies.

Configure OSPF as the IGP

Before you can run MPLS on your network, you must have an IGP running on all specified routers and interfaces. The IGP can be either OSPF or IS-IS. For the steps to configure IS-IS, see [“Configure IS-IS as the IGP” on page 2185](#).

Figure 165: OSPF Network Topology



The OSPF IGP in the MPLS network in [Figure 164 on page 2186](#) consists of the following:

- All routers are configured for the backbone OSPF area 0.0.0.0.
- All routers have the RID manually configured to avoid possible problems when the OSPF RID changes; for example, when multiple loopback addresses are configured.
- All routers have traffic engineering enabled. When traffic engineering is enabled for OSPF, the SPF algorithm takes into account the various LSPs configured under MPLS and configures OSPF to generate

link-state advertisements (LSAs) that carry traffic engineering parameters. These routes are installed into the primary routing table **inet.0**, but the LSPs are installed by default into the **inet.3** routing table.

- Adjacencies between all OSPF neighbors are as follows:

- R1 is adjacent to R2, R3, and R5
- R2 is adjacent to R1, R3, R4, and R6
- R3 is adjacent to R1, R2, R4, and R6
- R4 is adjacent to R2, R3, R5, and R6
- R5 is adjacent to R1, R4, and R6
- R6 is adjacent to R2, R3, R4, and R5

When you configure OSPF as the IGP, you must enable OSPF and traffic engineering on the router. We also recommend that you manually configure the RID and include the loopback interface (**lo0**) at the **[edit protocols ospf]** hierarchy level.

You can enable OSPF throughout the rest of the network by repeating this step as appropriate on other routers until all routers and interfaces establish OSPF neighbors.

To configure OSPF and establish OSPF neighbors, follow these steps:

1. [Enable OSPF on Routers in Your Network | 2197](#)
2. [Verify That OSPF Neighbors Are Established | 2200](#)

Enable OSPF on Routers in Your Network

Action

To enable OSPF on routers in your MPLS network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols ospf
```

2. Configure the area and the interface:

```
[edit protocols ospf]
user@host# set area area-id interface type-fpc/pic/port
```

3. Disable the management interface if you have included the **interface all** statement in the previous step:

```
[edit protocols ospf]
```



```
user@host# set interface fxp0.0 disable
```

4. Include the loopback (**lo0**) interface if you intend to manually configure the RID:

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface lo0
```

5. Set the loopback interface (**lo0**) to passive:

```
[edit protocols ospf]
user@host# set area 0.0.0.0 interface lo0 passive
```

6. Configure traffic engineering:

```
[edit protocols ospf]
user@host# set traffic-engineering
```

7. Manually configure the RID at the [**routing-options**] hierarchy level:

```
[edit]
user@host# edit routing-options
[edit routing-options]
user@host# set router-id router-id
```

8. Verify and commit the entire configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols ospf

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface so-0/0/0.0
```



```

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface lo0

[edit protocols ospf]
user@R6# set area 0.0.0.0 interface lo0 passive

[edit protocols ospf]
user@R6# set traffic-engineering

[edit protocols ospf]
user@R6# show
traffic-engineering;
area 0.0.0.0 {
    interface so-0/0/0.0;
    interface so-0/0/1.0;
    interface so-0/0/2.0;
    interface so-0/0/3.0;
    interface lo0.0 {
        passive;
    }
}

[edit protocols ospf]
user@R6# commit
commit complete

[edit]
user@R6# edit routing-options

[edit routing-options]
user@R6# set router-id 10.0.0.6

[edit routing-options]
user@R6# show
[...Output truncated...]
router-id 10.0.0.6;
autonomous-system 65432;

[edit routing-options]
user@R6# commit
commit complete

```


Meaning

The sample output shows that OSPF, with traffic engineering, is enabled on the interfaces on egress router **R6**. In addition, the RID is configured manually to avoid possible problems when the OSPF RID changes; for example, when multiple loopback addresses are configured. The RID uniquely identifies the router within the OSPF network. It is transmitted within the LSAs used to populate the link-state database and calculate the shortest-path tree. In a link-state network, it is important that two routers do not share the same RID value, otherwise IP routing problems may occur.

The sample outputs also shows the **passive** statement included in the configuration of the loopback (**lo0**) interface. Including the **passive** statement is considered best practice and ensures the following:

- Protocols are not run over the loopback (**lo0**) interface
- When the router ID (RID) is configured manually, ensures that the loopback (**lo0**) interface is advertised to other networks.

Verify That OSPF Neighbors Are Established

Purpose

After configuring OSPF, you must verify that neighboring routers have formed adjacencies with each other.

Action

To verify OSPF neighbors, enter the following Junos OS CLI operational mode command:

```
user@host> show ospf neighbor
```

Sample Output

```
user@R1> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.12.2	so-0/0/0.0	Full	10.0.0.2	128	37
10.1.15.2	so-0/0/1.0	Full	10.0.0.5	128	35
10.1.13.2	so-0/0/2.0	Full	10.0.0.3	128	38

```
user@R3> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.34.2	so-0/0/0.0	Full	10.0.0.4	128	38
10.1.23.1	so-0/0/1.0	Full	10.0.0.2	128	35
10.1.13.1	so-0/0/2.0	Full	10.0.0.1	128	37
10.1.36.2	so-0/0/3.0	Full	10.0.0.6	128	36


```
user@R6> show ospf neighbor
```

Address	Interface	State	ID	Pri	Dead
10.1.56.1	so-0/0/0.0	Full	10.0.0.5	128	39
10.1.46.1	so-0/0/1.0	Full	10.0.0.4	128	37
10.1.26.1	so-0/0/2.0	Full	10.0.0.2	128	36
10.1.36.1	so-0/0/3.0	Full	10.0.0.3	128	37

Meaning

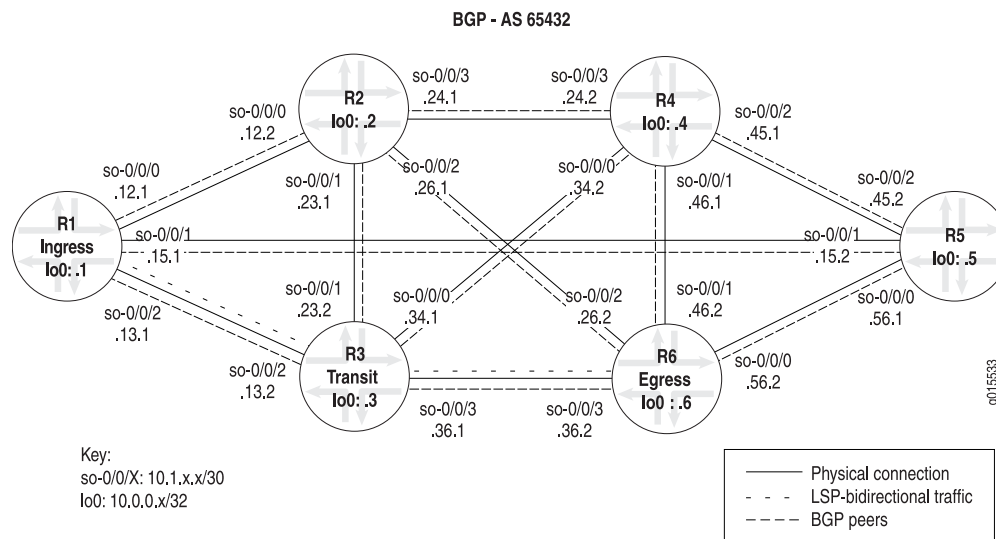
The sample output from the ingress, transit, and egress routers shows that all routers in the network have formed OSPF neighbor adjacencies.

Set Up BGP on Routers in Your Network

Before BGP can function in your MPLS network, you must define the autonomous system (AS) number on the routers in your network, and configure at least one group that includes at least one peer.

Optionally, you can configure a routing policy. The routing policy allows you to control the information shared with BGP neighbors and provides the opportunity to filter and modify the information you receive.

Figure 166: BGP Network Topology



The BGP configuration in the MPLS network in [Figure 166 on page 2201](#) consists of the following:

- A full-mesh IBGP topology, using AS 65432.
- All IBGP sessions peer between loopback addresses because significant stability advantages are gained.

- All routers are configured with one group, **group internal**.
- A **send-statics** policy on routers R1 and R6 allows a new route to be advertised into the network.

The example network uses IS-IS Level 2 and a policy to create routes that are reachable through the LSP. However, IS-IS Level 1 or an OSPF area can be used and the policy omitted if the network has existing BGP traffic.

You can set up BGP throughout the rest of the network by repeating Step 1, [“Define the Local Autonomous System” on page 2202](#) through Step 3, [“Configure a Simple Routing Policy” on page 2205](#) as appropriate on other routers until all routers are set up with BGP.

To set up BGP on routers in your network, follow these steps:

1. [Define the Local Autonomous System | 2202](#)
2. [Configure BGP Neighbor Connections | 2203](#)
3. [Configure a Simple Routing Policy | 2205](#)
4. [Verify That BGP Sessions Are Up | 2208](#)

Define the Local Autonomous System

Purpose

Before BGP can function, you need to define a local AS number on the routers in your network. In the example network in [Figure 166 on page 2201](#), all routers are in AS 65432.

Action

To define an AS number on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit routing-options
```

2. Configure all interfaces to a specific AS:

```
[edit routing-options]
user@host# set autonomous-system as-number
```

3. Verify the configuration:

```
user@host# show
```

```
user@host# commit
```


Sample Output

```

user@R1> edit
Entering configuration mode

[edit]
user@R1# edit routing-options

[edit routing-options]
user@R1# set autonomous-system 65432

[edit routing-options]
user@R1# show
[...Output truncated...]
autonomous-system 65432;

[edit routing-options]
user@R6# commit
commit complete

```

Meaning

The output shows that router **R1** resides in **AS 65432**. All other routers in the example network shown in [Figure 166 on page 2201](#) also reside in **AS 65432**.

Configure BGP Neighbor Connections

Purpose

You must configure at least one group that includes at least one peer for BGP to run in your network. First determine which neighbors are internal or external to your local AS boundary. Internal neighbors are inside your local AS boundary. In the example network shown in [Figure 166 on page 2201](#), all the routers are in one AS and are therefore internal. In this example, all IBGP sessions peer between loopback addresses because significant stability advantages are gained. For more information about configuring BGP neighbor connections, see the *Junos Routing Protocols Configuration Guide*.

Action

To configure BGP neighbor connections, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```

[edit]
user@host# edit protocols bgp

```

2. Configure the group and peer's IP address:


```
[edit protocols bgp]
user@host# set group group-name type type neighbor neighbor-address
```

NOTE: For external neighbors, use the following form of the command that includes the peer's AS number:

```
user@host# set group group-name neighbor neighbor-address peer-as peer-as-number
```

3. Configure the local address:

```
[edit protocols bgp]
user@host# set group group-name local-address local-address
```

4. Verify and commit the configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols bgp

[edit protocols bgp]
user@R1# set group internal type internal neighbor 10.0.0.2

[edit protocols bgp]
users@R1# set group internal local-address 10.0.0.1

[edit protocols bgp]
user@R1# show
group internal {
    type internal;
    local-address 10.0.0.1;
```



```

    neighbor 10.0.0.2;
    neighbor 10.0.0.3;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
}

[edit protocols bgp]
user@R1# commit
commit complete

```

Meaning

The sample output shows that router **R1** is in an internal group with five BGP neighbors. The **local-address** statement is included in this example configuration because IBGP is used. It is considered best practice to configure a local address when you use an IBGP. BGP messages are sourced from the loopback address because the **local-address** statement is included in the configuration. Generally, you would not configure a local address when external BGP is configured.

Configure a Simple Routing Policy

Purpose

Routing policy allows you to control the information shared with BGP neighbors and provides the opportunity to filter and modify the information you receive. Typically, a network is injected into BGP using a policy. This may also be done through a static route. In the network in [Figure 166 on page 2201](#), a static route export policy is used to inject routes into BGP.

Action

To configure a simple routing policy, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```

[edit]
user@host# edit routing-options

```

2. Configure a static route for redistribution to other autonomous systems:

```

[edit routing-options]
user@host# set static route destination/24 reject

```

3. Configure a routing policy that matches and accepts the configured static routes into BGP updates:


```
[edit]
user@host# edit policy-options
[edit policy-options]
user@host# set policy-statement policy-name term term-name from route-filter address exact
user@host# set policy-statement policy-name term term-name then accept
```

4. Apply the policy created in Step 3 to all BGP neighbors:

```
[edit]
user@host# edit protocols bgp
[edit protocols bgp]
user@host# set export policy-name
```

5. Verify and commit the configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit routing-options

[edit routing-options]
user@R1# set static route 100.100.1.0/24 reject

[edit routing-options]
user@R1# show
[...Output truncated...]
    route 100.100.1.0/24 reject;
}
router-id 10.0.0.1;
autonomous-system 65432;

[edit routing-options]
user@R1# top
```



```

[edit]
user@R1# edit policy-options

[edit policy-options]
user@R1# set policy-statement send-statics term statics from route-filter 100.100.1.0/24 exact

[edit policy-options]
user@R1# set policy-statement send-statics term statics then accept

[edit policy-options]
user@R1# top

[edit]
user@R1# edit protocols bgp

[edit protocols bgp]
user@R1# set export send-statics

[edit protocols bgp]
user@R1# show
export send-statics;
group internal {
    type internal;
    local-address 10.0.0.1;
    neighbor 10.0.0.2;
    neighbor 10.0.0.3;
    neighbor 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.6;
}

[edit protocols bgp]
user@R1# commit
commit complete

```

Meaning

The sample output shows that routing policy **send-statics** is configured on the router. The routing policy matches and accepts the configured static routes into the routing table and injects the routes into BGP updates. Typically, a routing policy is applied at the group level, although it can be applied at the global level, as shown in this example.

Verify That BGP Sessions Are Up

Purpose

After configuring BGP, you must verify that BGP peers are established and the sessions are up.

Action

To verify BGP peers and sessions, enter the following Junos OS CLI operational mode command:

```
user@host> show bgp summary
```

Sample Output

```
user@R1> show bgp summary
```

```
Groups: 1 Peers: 5 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	1	1	0	0	0	0	0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last Up/Dwn
State #Active/Received/Damped...						
10.0.0.2	65432	1369	1373	0	0	11:25:11 0/0/0
	0/0/0					
10.0.0.3	65432	1369	1372	0	0	11:24:55 0/0/0
	0/0/0					
10.0.0.4	65432	1369	1372	0	0	11:25:03 0/0/0
	0/0/0					
10.0.0.5	65432	1369	1372	0	0	11:25:07 0/0/0
	0/0/0					
10.0.0.6	65432	1343	1344	0	1	11:10:55 1/1/0
	0/0/0					

```
user@R3> show bgp summary
```

```
Groups: 1 Peers: 4 Down peers: 0
```

Table	Tot Paths	Act Paths	Suppressed	History	Damp	State	Pending
inet.0	2	2	0	0	0	0	0

Peer	AS	InPkt	OutPkt	OutQ	Flaps	Last Up/Dwn
State #Active/Received/Damped...						
10.0.0.1	65432	1375	1375	0	6	11:26:57 1/1/0
	0/0/0					
10.0.0.2	65432	43016	43016	0	0	2w0d22h 0/0/0
	0/0/0					
10.0.0.4	65432	74460	74461	0	0	3w4d20h 0/0/0
	0/0/0					


```

10.0.0.6          65432          1347          1347          0          6      11:13:10 1/1/0
                  0/0/0

user@R6> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State   Pending
inet.0                1          1          0          0          0          0
Peer            AS        InPkt      OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.1        65432          1348          1350        0        0    11:13:46 1/1/0
                  0/0/0
10.0.0.2        65432          1347          1351        0        0    11:14:02 0/0/0
                  0/0/0
10.0.0.3        65432          1347          1350        0        0    11:13:58 0/0/0
                  0/0/0
10.0.0.4        65432          1347          1350        0        0    11:13:54 0/0/0
                  0/0/0
10.0.0.5        65432          1347          1350        0        0    11:13:50 0/0/0
                  0/0/0

```

Meaning

The sample output from the ingress, transit, and egress routers shows that all routers in the network shown in [Figure 166 on page 2201](#) have BGP peers established and sessions up.

Define the Local Autonomous System

Purpose

Before BGP can function, you need to define a local AS number on the routers in your network. In the example network in [Figure 166 on page 2201](#), all routers are in AS 65432.

Action

To define an AS number on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```

[edit]
user@host# edit routing-options

```

2. Configure all interfaces to a specific AS:


```
[edit routing-options]  
user@host# set autonomous-system as-number
```

3. Verify the configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R1> edit  
Entering configuration mode  
  
[edit]  
user@R1# edit routing-options  
  
[edit routing-options]  
user@R1# set autonomous-system 65432  
  
[edit routing-options]  
user@R1# show  
[...Output truncated...]  
autonomous-system 65432;  
  
[edit routing-options]  
user@R6# commit  
commit complete
```

Meaning

The output shows that router **R1** resides in **AS 65432**. All other routers in the example network shown in [Figure 166 on page 2201](#) also reside in **AS 65432**.

Enable MPLS and RSVP

You can enable MPLS and RSVP throughout the rest of the network by repeating Step 1, “[Enable MPLS and RSVP on Routers](#)” on page 2211 and Step 2, “[Enable MPLS on Transit Router Interfaces](#)” on page 2213 as appropriate on other routers until all routers are enabled with MPLS and RSVP.

1. [Enable MPLS and RSVP on Routers](#) | 2211
2. [Enable MPLS on Transit Router Interfaces](#) | 2213

Enable MPLS and RSVP on Routers

Action

To enable MPLS and RSVP on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols
```

2. Configure MPLS and RSVP:

```
[edit protocols]
user@host# set mpls interface all
user@host# set rsvp interface all
```

3. Disable the management interface for MPLS and RSVP:

```
[edit protocols mpls]
user@host# set interface fxp0.0 disable
[edit protocols rsvp]
user@host# set interface fxp0.0 disable
```

4. Verify and commit the configuration:

```
user@host# show
```

```
user@host# commit
```


Sample Output

```

user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols

[edit protocols]
user@R1# set mpls interface all

[edit protocols]
user@R1# set rsvp interface all

[edit protocols]
user@R1# show
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

[edit protocols]
user@R1# commit
commit complete

```

Meaning

The sample output shows that router **R1** has MPLS and RSVP enabled on all interfaces, except for the management interface (**fxp0.0**), which is disabled. It is considered best practice to disable the management interlace (**fxp0.0**) for MPLS and RSVP to preempt any problems.

Typically every interface that you want to use is listed. For an example of a router configured with specific interfaces, see [“Enable IS-IS on Routers in Your Network” on page 2187](#).

Enable MPLS on Transit Router Interfaces

Purpose

Even though transit interfaces are enabled with MPLS when you include the **family mpls** statement in the configuration, MPLS as a whole is not configured on your router or in your network. You must complete all five steps in this topic to have the MPLS protocol running on your network.

NOTE: The management interface (**fxp0**) and the loopback interface (**lo0**) are not transit interfaces.

Action

To configure transit interfaces to support MPLS, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure MPLS:

```
[edit interfaces]
user@host# set type-fpc/pic/port unit number family mpls
```

3. Verify and commit the configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set so-0/0/2 unit 0 family mpls
```



```
[edit interfaces]
user@R1# show
so-0/0/2 {
  unit 0 {
    family inet {
      address 10.1.13.1/30;
    }
    family iso;
    family mpls;
  }
}

[edit interfaces]
user@R1# commit
commit complete
```

Meaning

The sample output shows that the interface **so-0/0/2** is configured to support MPLS. The **family** statement identifies which protocol packets are accepted into the interfaces. For example, valid MPLS packets are dropped if the interface is not configured with the MPLS protocol.

Enable MPLS and RSVP on Routers

Action

To enable MPLS and RSVP on routers in your network, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols
```

2. Configure MPLS and RSVP:

```
[edit protocols]
user@host# set mpls interface all
user@host# set rsvp interface all
```

3. Disable the management interface for MPLS and RSVP:


```
[edit protocols mpls]
user@host# set interface fxp0.0 disable
[edit protocols rsvp]
user@host# set interface fxp0.0 disable
```

4. Verify and commit the configuration:

```
user@host# show
```

```
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit protocols

[edit protocols]
user@R1# set mpls interface all

[edit protocols]
user@R1# set rsvp interface all

[edit protocols]
user@R1# show
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

[edit protocols]
```



```
user@R1# commit
commit complete
```

Meaning

The sample output shows that router **R1** has MPLS and RSVP enabled on all interfaces, except for the management interface (**fxp0.0**), which is disabled. It is considered best practice to disable the management interface (**fxp0.0**) for MPLS and RSVP to preempt any problems.

Typically every interface that you want to use is listed. For an example of a router configured with specific interfaces, see [“Enable IS-IS on Routers in Your Network” on page 2187](#).

Enable MPLS on Transit Router Interfaces

Purpose

Even though transit interfaces are enabled with MPLS when you include the **family mpls** statement in the configuration, MPLS as a whole is not configured on your router or in your network. You must complete all five steps in this topic to have the MPLS protocol running on your network.

NOTE: The management interface (**fxp0**) and the loopback interface (**lo0**) are not transit interfaces.

Action

To configure transit interfaces to support MPLS, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit interfaces
```

2. Configure MPLS:

```
[edit interfaces]
user@host# set type-fpc/pic/port unit number family mpls
```

3. Verify and commit the configuration:

```
user@host# show
```



```
user@host# commit
```

Sample Output

```
user@R1> edit
Entering configuration mode

[edit]
user@R1# edit interfaces

[edit interfaces]
user@R1# set so-0/0/2 unit 0 family mpls

[edit interfaces]
user@R1# show
so-0/0/2 {
  unit 0 {
    family inet {
      address 10.1.13.1/30;
    }
    family iso;
    family mpls;
  }
}

[edit interfaces]
user@R1# commit
commit complete
```

Meaning

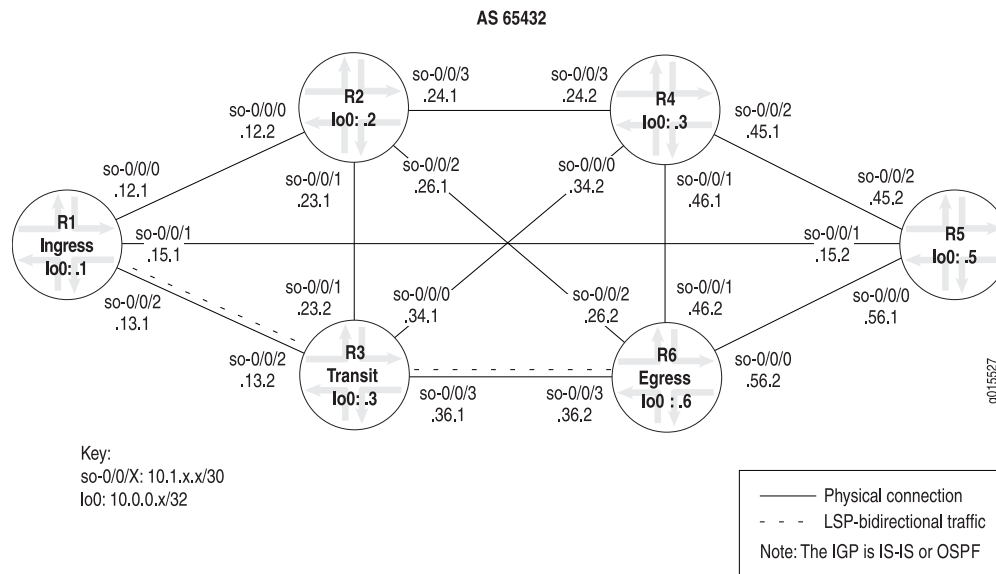
The sample output shows that the interface **so-0/0/2** is configured to support MPLS. The **family** statement identifies which protocol packets are accepted into the interfaces. For example, valid MPLS packets are dropped if the interface is not configured with the MPLS protocol.

Verifying the MPLS Configuration

After configuring MPLS on your network, you must verify the correct configuration of both the MPLS and RSVP protocols. Incorrect configuration of either protocol prevents successful LSP creation.

Figure 167 on page 2218 illustrates the network with the example configurations used in this topic.

Figure 167: MPLS Network Topology



To verify the MPLS configuration, follow these steps:

1. [Verify the RSVP Protocol | 2218](#)
2. [Verify RSVP Interfaces | 2220](#)
3. [Verify Protocol Families | 2222](#)

Verify the RSVP Protocol

Purpose

If the RSVP protocol is not enabled on the routers in your network, the interface cannot signal LSPs.

Action

To verify that the RSVP protocol is enabled, enter the following Junos OS CLI command:

```
user@host> show rsvp version
```

Sample Output

```
user@R1> 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 restart        = Disabled
Restart helper mode     = Enabled
Restart time            = 0 msec

```

Meaning

The sample output shows that the RSVP protocol is enabled on **R1**. The supported RSVP protocol is version 1, as defined in RFC 2205.

The RSVP refresh timer is set to 30 seconds, indicating that every 30 seconds, plus or minus 50 percent, the router will refresh the RSVP state with its directly connected neighbors by sending either a **Path** or a **Resv** message. The variable refresh time helps prevent harmonic oscillations in network traffic caused by periodic protocol updates.

The keepalive multiplier, **K(keep multiplier)**, is input to a formula that helps determine the lifetime of an RSVP session. The session lifetime is reset each time the state is updated. The lifetime represents the duration of an RSVP session that does not receive any state updates (**Path** or **Resv** messages). The formula is:

$$\text{RSVP session lifetime} = (\text{keep-multiplier} + 0.5) * 1.5 * \text{refresh-time}$$

The RSVP **preemption** state is currently configured for normal preemption, indicating that only an LSP with a stronger priority can preempt an existing session; that is, the setup value of the new LSP is lower than the hold value of the existing LSP. Other options include **aggressive** preemption, which always preempts when there is insufficient bandwidth, and **disabled**, which prevents any preemption, regardless of LSP priority values.

Graceful restart is currently disabled and **Restart helper mode** is enabled. There are four combinations for **Graceful restart** and **restart helper mode**:

1. Both **Graceful restart** and **Restart helper mode** are enabled.
2. **Graceful restart** is enabled but **Restart helper mode** is disabled. An LSR with this configuration can restart gracefully but cannot help a neighbor with its restart and recovery procedures.
3. **Graceful restart** is disabled but **Restart helper mode** is enabled. An LSR with this configuration can only help a restarting neighbor. It cannot restart gracefully itself.
4. **Graceful restart** and **Restart helper mode** are both disabled. This configuration completely disables RSVP graceful restart (including restart and recovery procedures and helper mode). It is the same as an LSR that is not supported by RSVP graceful restart.

Restart time is the estimated time (in milliseconds) for an LSR to restart the RSVP traffic engineering component. In the example output, the restart time is 0 milliseconds, indicating that it is disabled.

The output is identical for all routers in the network shown in [“MPLS Network Topology” on page 2217](#).

Verify RSVP Interfaces

Purpose

If the RSVP protocol is not configured correctly on the routers in your network, the interfaces cannot signal LSPs.

Action

To verify RSVP interfaces, enter the following Junos OS CLI operational mode command:

```
user@host> show rsvp interface
```

Sample Output 1

```
user@R1> show rsvp interface
```

```
RSVP interface: 4 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	2	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R2> show rsvp interface
```

```
RSVP interface: 5 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

```
user@R3> show rsvp interface
```

```
RSVP interface: 5 active
```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps


```
so-0/0/2.0 Up          0 100% 155.52Mbps 155.52Mbps 0bps      0bps
so-0/0/3.0 Up          0 100% 155.52Mbps 155.52Mbps 0bps      0bps
```

user@R4> **show rsvp interface**

RSPV interface: 5 active

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

user@R5> **show rsvp interface**

RSPV interface: 4 active

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

user@R6> **show rsvp interface**

RSPV interface: 5 active

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/3.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

Sample Output 2

user@R6> **show rsvp interface**

RSPV interface: 3 active

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
so-0/0/0.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/1.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps
so-0/0/2.0	Up	0	100%	155.52Mbps	155.52Mbps	0bps	0bps

#so-0/0/3.0 is missing

Sample Output 3

```
user@host# show rsvp interface
RSVP not configured
```

Meaning

Sample Output 1 shows that all interfaces on all routers in the network are enabled with RSVP, including the management interface (**fxp0**). The output for all routers in the network includes similar information, so we will examine R6 in detail.

R6 has five interfaces enabled with RSVP (**Up**). Interface **so-0/1/1.0** has a single active RSVP reservation (**Active resv**) that did not change the default subscription percentage of 100 percent (**Subscription**).

Interface **so-0/1/1.0** did not assign a static bandwidth (**Static BW**) to the logical unit and therefore inherited 100 percent of the physical interface rate as the bandwidth available (**Available BW**) for RSVP sessions. Interface **so-0/1/1.0** has no bandwidth assigned (**Reserved BW**), and no RSVP bandwidth allocation at any single instant in time (**Highwater mark**).

Sample Output 2 shows that interface **so-0/0/3.0** is missing. If you do not configure the correct interface at the **[edit protocols rsvp]** hierarchy level, the interface cannot signal LSPs, and does not appear in the output for the **show rsvp interface** command.

Sample Output 3 shows that the RSVP protocol is not configured at the **[edit protocols rsvp]** hierarchy level.

Verify Protocol Families

Purpose

If a logical interface does not have MPLS enabled, it cannot perform MPLS switching. This step allows you to quickly determine which interfaces are configured with MPLS and other protocol families.

Action

To verify the protocol families configured on the routers in your network, enter the following Junos OS CLI operational mode command:

```
user@host> show interfaces terse
```


Sample Output 1

user@R1> show interfaces terse

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.12.1/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.15.1/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.13.1/30	
			iso		
			mpls		
so-0/0/3	up	down			

user@R2> show interfaces terse

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.12.2/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.23.1/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.26.1/30	
			iso		
			mpls		
so-0/0/3	up	up			
so-0/0/3.0	up	up	inet	10.1.24.1/30	
			iso		
			mpls		

user@R3> show interfaces terse

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.34.1/30	
			iso		
			mpls		


```

so-0/0/1          up    up
so-0/0/1.0        up    up    inet  10.1.23.2/30
                    iso
                    mpls

so-0/0/2          up    up
so-0/0/2.0        up    up    inet  10.1.13.2/30
                    iso
                    mpls

so-0/0/3          up    up
so-0/0/3.0        up    up    inet  10.1.36.1/30
                    iso
                    mpls

```

user@R4> **show interfaces terse**

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.34.2/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.46.1/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.45.1/30	
			iso		
			mpls		
so-0/0/3	up	up			
so-0/0/3.0	up	up	inet	10.1.24.2/30	
			iso		
			mpls		

user@R5> **show interfaces terse**

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.56.1/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.15.2/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.45.2/30	


```

iso
mpls
so-0/0/3          up    down

user@R6> show interfaces terse
Interface          Admin Link Proto Local          Remote
so-0/0/0           up    up
so-0/0/0.0         up    up   inet  10.1.56.2/30
                  iso
                  mpls
so-0/0/1           up    up
so-0/0/1.0         up    up   inet  10.1.46.2/30
                  iso
                  mpls
so-0/0/2           up    up
so-0/0/2.0         up    up   inet  10.1.26.2/30
                  iso
                  mpls
so-0/0/3           up    up
so-0/0/3.0         up    up   inet  10.1.36.2/30
                  iso
                  mpls

```

Sample Output 2

```

user@R6> show interfaces terse
Interface          Admin Link Proto Local          Remote
so-0/0/0           up    up
so-0/0/0.0         up    up   inet  10.1.56.2/30
                  iso
                  mpls
so-0/0/1           up    up
so-0/0/1.0         up    up   inet  10.1.46.2/30
                  iso
                  mpls
so-0/0/2           up    up
so-0/0/2.0         up up inet 10.1.26.2/30
                  iso #The mpls statement is missing.
so-0/0/3           up    up
so-0/0/3.0         up    up   inet  10.1.36.2/30

```



```
iso
mpls
```

Meaning

Sample Output 1 shows the interface, the administrative status of the link (**Admin**), the data link layer status of the link (**Link**), the protocol families configured on the interface (**Proto**), and the local and remote addresses on the interface.

All interfaces on all routes in the network shown in [Figure 167 on page 2218](#) are administratively enabled and functioning at the data link layer with MPLS and IS-IS, and have an **inet** address. All are configured with an IPv4 protocol family (**inet**), and have the IS-IS (**iso**) and MPLS (**mpls**) protocol families configured at the `[edit interfaces type-fpc/pic/port unit number]` hierarchy level.

Sample Output 2 shows that interface **so-0/0/2.0** on **R6** does not have the **mpls** statement included at the `[edit interfaces type-fpc/pic/port unit number]` hierarchy level.

Verify the RSVP Protocol

Purpose

If the RSVP protocol is not enabled on the routers in your network, the interface cannot signal LSPs.

Action

To verify that the RSVP protocol is enabled, enter the following Junos OS CLI command:

```
user@host> show rsvp version
```

Sample Output

```
user@R1> 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 restart        = Disabled
```



```
Restart helper mode    = Enabled
Restart time           = 0 msec
```

Meaning

The sample output shows that the RSVP protocol is enabled on **R1**. The supported RSVP protocol is version 1, as defined in RFC 2205.

The RSVP refresh timer is set to 30 seconds, indicating that every 30 seconds, plus or minus 50 percent, the router will refresh the RSVP state with its directly connected neighbors by sending either a **Path** or a **Resv** message. The variable refresh time helps prevent harmonic oscillations in network traffic caused by periodic protocol updates.

The keepalive multiplier, **K(keep multiplier)**, is input to a formula that helps determine the lifetime of an RSVP session. The session lifetime is reset each time the state is updated. The lifetime represents the duration of an RSVP session that does not receive any state updates (**Path** or **Resv** messages). The formula is:

$$\text{RSVP session lifetime} = (\text{keep-multiplier} + 0.5) * 1.5 * \text{refresh-time}$$

The RSVP **preemption** state is currently configured for normal preemption, indicating that only an LSP with a stronger priority can preempt an existing session; that is, the setup value of the new LSP is lower than the hold value of the existing LSP. Other options include **aggressive** preemption, which always preempts when there is insufficient bandwidth, and **disabled**, which prevents any preemption, regardless of LSP priority values.

Graceful restart is currently disabled and **Restart helper mode** is enabled. There are four combinations for **Graceful restart** and **restart helper mode**:

1. Both **Graceful restart** and **Restart helper mode** are enabled.
2. **Graceful restart** is enabled but **Restart helper mode** is disabled. An LSR with this configuration can restart gracefully but cannot help a neighbor with its restart and recovery procedures.
3. **Graceful restart** is disabled but **Restart helper mode** is enabled. An LSR with this configuration can only help a restarting neighbor. It cannot restart gracefully itself.
4. **Graceful restart** and **Restart helper mode** are both disabled. This configuration completely disables RSVP graceful restart (including restart and recovery procedures and helper mode). It is the same as an LSR that is not supported by RSVP graceful restart.

Restart time is the estimated time (in milliseconds) for an LSR to restart the RSVP traffic engineering component. In the example output, the restart time is 0 milliseconds, indicating that it is disabled.

The output is identical for all routers in the network shown in [“MPLS Network Topology” on page 2217](#).

Define a Load-Balancing Policy

Purpose

On the ingress or transit router, you can include a policy statement that performs load balancing on all routes. For information on including a policy statement that performs load balancing on specific routes, see “Configuring Per-Packet Load Balancing” in the *Junos Routing Protocols Configuration Guide*.

Action

On the ingress or transit router, to define a load-balancing policy for all routes, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit policy-options
```

2. Define the load-balance policy and action:

```
[edit policy-options]
user@host# set policy-statement policy-name then load-balance per-packet
```

3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit policy-options

[edit policy-options]
user@R6# set policy-statement load-balance-traffic then load-balance per-packet

[edit policy-options]
user@R6# show
policy-statement load-balance-traffic {
    then {
        load-balance per-packet;
    }
}
```



```

}

[edit policy-options]
user@R6# commit
commit complete

```

Meaning

The sample output from ingress router **R6** shows the process for configuring load balancing. On an Internet Processor I ASIC, packets with the same parameters are spread across multiple equal-cost next hops; while an Internet Processor II ASIC sends packets with the same parameters to the same next hop, since they are in the same flow. The Junos OS command to turn on load balancing uses the action **load-balance per-packet**, which is misnamed in relation to the Internet Processor II ASIC. On the Internet Processor II ASIC, this command actually enables per-flow load balancing.

Use the traceroute Command to Verify MPLS Labels

Purpose

You can use the **traceroute** command to verify that packets are being sent over the LSP.

Action

To verify MPLS labels, enter the following Junos OS CLI operational mode command, where **host-name** is the IP address or the name of the remote host:

```
user@host> traceroute host-name
```

Sample Output 1

```

user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.12.2 (10.1.12.2)  0.861 ms  0.718 ms  0.679 ms
    MPLS Label=100048 CoS=0 TTL=1 S=1
 2  10.1.24.2 (10.1.24.2)  0.822 ms  0.731 ms  0.708 ms
    MPLS Label=100016 CoS=0 TTL=1 S=1
 3  10.1.46.2 (10.1.46.2)  0.571 ms !N  0.547 ms !N  0.532 ms !N

```


Sample Output 2

```
user@R1> traceroute 10.0.0.6
traceroute to 10.0.0.6 (10.0.0.6), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.605 ms  0.548 ms  0.503 ms
 2  10.0.0.6 (10.0.0.6)  0.761 ms  0.676 ms  0.675 ms
```

Meaning

Sample Output 1 shows that MPLS labels are used to forward packets through the network. Included in the output is a label value (**MPLS Label=100048**), the time-to-live value (**TTL=1**), and the stack bit value (**S=1**).

The **MPLS Label** field is used to identify the packet to a particular LSP. It is a 20-bit field, with a maximum value of $(2^{20}-1)$, or approximately 1,000,000.

The TTL value contains a limit on the number of hops that this MPLS packet can travel through the network (1). It is decremented at each hop, and if the TTL value drops below one, the packet is discarded.

The bottom of the stack bit value (**S=1**) indicates that is the last label in the stack and that this MPLS packet has one label associated with it. The MPLS implementation in the Junos OS supports a stacking depth of 3 on the M-series routers and up to 5 on the T-series platforms. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

MPLS labels appear in Sample Output 1 because the **traceroute** command is issued to a BGP destination where the BGP next hop for that route is the LSP egress address. The Junos OS default behavior uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Sample Output 2 shows that MPLS labels do not appear in the output for the **traceroute** command. If the BGP next hop does not equal the LSP egress address or the destination is an IGP route, the BGP traffic does not use the LSP. Instead of using the LSP, the BGP traffic is using the IGP (IS-IS, in this case) to reach the egress address (**R6**).

Apply the Load-Balancing Policy to the Forwarding Table

Purpose

Apply the policy configured in Step 1 to routes exported from the routing table to the forwarding table.

Action

To apply a load-balancing policy to the forwarding table, follow these steps:

1. In configuration mode, go to the following hierarchy level:


```
[edit]
user@host# edit routing-options
```

2. Define a load-balance per packet action:

```
[edit routing-options]
user@host# set forwarding-table export policy-name
```

3. Verify and commit the configuration:

```
user@host# show
user@host# commit
```

Sample Output

```
[edit]
user@R6# edit routing-options

[edit routing-options]
user@R6# set forwarding-table export load-balance-traffic

[edit routing-options]
user@R6# show
static {
[...Output truncated...]
}
router-id 192.168.6.1;
autonomous-system 65432;
forwarding-table {
    export load-balance-traffic;
}

[edit routing-options]
user@R6# commit
commit complete
```

Meaning

The sample output shows the process for applying a load-balancing policy to export routes from the routing table to the forwarding table.

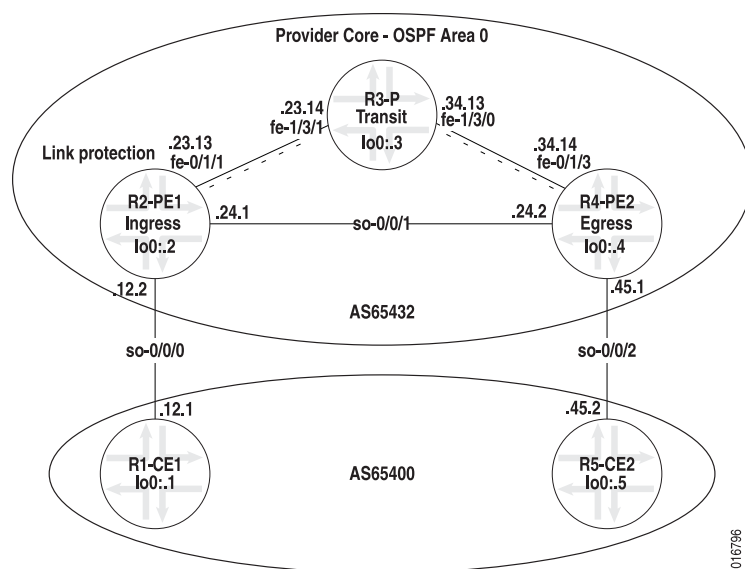
Fast Reroute Problem Overview

Problem

Description: Incorrect configuration is a common mistake when trying to establish protection for an MPLS LSP. Protection with either fast reroute or link protection requires a **per-packet load-balance policy** exported at the [edit routing-options forwarding-table] hierarchy level. Correctly configured protection for an MPLS LSP results in two next-hop forwarding table entries per destination, either an incoming MPLS label or an IP destination.

Figure 168 on page 2232 illustrates a network topology with link protection and load balancing enabled to ensure that routes are correctly placed in to the forwarding table.

Figure 168: Fast Reroute Problem Network



The network shown in Figure 168 on page 2232 illustrates an MPLS-based VPN with traffic protection and load balancing, consisting of the following:

- All physical interfaces addresses are from the 10.0.x.x/30 address space.
- All loopback addresses are from the 192.168.x.1/32 block.
- The IGP is a single-area (Area 0) OSPF.
- RSVP is deployed as the MPLS signaling protocol between PE routers.
- LSPs (r2-r4 and r4-r2) established between PE routers.
- MP-IBGP mesh between PE routers, loopback peering, and VPN-IPv4 NLRI.
- CE-PE link running EBGp.
- Full-mesh Layer 3 VPN between CE1 and CE2.

- Traffic protection for the link between the PE1 and P routers.
- Load balancing on PE1.

The overall goal of this network is to provide point-to-point connectivity between the two CE routers and traffic protection in the core of the network.

Symptom

In the network shown in [Figure 168 on page 2232](#), the external symptom is that local repair is taking about one second to complete, which is slow. Use the **show route forwarding-table vpn vpn-a destination** command to check that the correct routes are included in the forwarding table. In the example output below, there is only one route installed in the forwarding table, when for fast local repair, there should be multiple next hops installed.

Sample Output

```
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive
Routing table: vpn-a.inet [Index 2]
Internet:

Destination: 192.168.5.0/24
Route type: user
Route reference: 0                               Route interface-index: 0
Flags: sent to PFE, prefix load balance
Next-hop type: indirect                          Index: 262142   Reference: 2
Next-hop type: Push 100160
Next-hop interface: so-0/0/1.0 #Only one next hop in the forwarding table.
```

Cause

Slow local repair is caused by the forwarding table not including the necessary next-hops to support local repair. The forwarding table shows only a single next-hop, when local repair requires additional next-hops for fast recovery.

Troubleshooting Commands

The Junos OS includes commands that are useful when troubleshooting a problem. This topic provides a brief description of each command followed by sample output, and a discussion of the output in relation to the problem.

The following commands can be used when troubleshooting a fast reroute error in an MPLS-VPN network:

```
user@R2-PE1> show configuration routing-instances vpn-a
user@R2-PE1> show configuration routing-options
user@R2-PE1> show bgp summary instance vpn-a
```



```

user@R2-PE1> show configuration protocols mpls
user@R2-PE1> show mpls lsp ingress
user@R2-PE1> show rsvp session ingress
user@R2-PE1> show rsvp session ingress detail
user@R2-PE1> show route table vpn-a 192.168.5.1 detail
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive

```

Sample Output

The **show configuration *statement-path*** command is used to display a specific configuration hierarchy; in this case, to verify the correct configuration of a specific routing instance named **vpn-a**.

```

user@R2-PE1> show configuration routing-instances vpn-a
instance-type vrf ;
interface so-0/0/0.0 ;
vrf-target {
    import target:65432:100;
    export target:65432:100;
}
protocols {
    bgp {
        group CE1 {
            type external;
            peer-as 65400;
            neighbor 10.0.12.1 ;
        }
    }
}

```

Meaning

The sample output for the **show configuration** command shows the current running configuration of the specific routing instance named **vpn-a** configured on the ingress PE1 router. The **vpn-a** instance configuration has a VRF table that supports EBGp routing on the PE-CE link (**so-0/0/0.0**). This interface is the correct interface for the CE1-PE1 link in the network topology shown in [Figure 168 on page 2232](#).

The VRF instance is linked to a VFR target community configured at the [edit policy-options] hierarchy level, allowing advertising of L3 VPN routes between PE routers. (See the PE1 configuration in [“Router Configurations” on page 2241](#) for the policy options configuration.) The import statement places, into the vpn-a.inet.0 table, all received L3 VPN MP-BGP routes tagged with the correct target community. The export statement advertises and tags all routes in the vpn-a.inet.o table with the listed target community to all MP-BGP peers.

The BGP protocols configuration within the routing instance applies the BGP import and export policies to the exchange of BGP routes on the PE-CE routing instance.

Sample Output

The **show bgp summary** command is used to display summary information about BGP and its neighbors to determine if routes are received from peers in the autonomous system (AS). In this case, information for the specified instance **vpn-a** is displayed.

```
user@R2-PE1> show bgp summary instance vpn-a
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
vpn-a.inet.0      11         7         0         0         0         0
Peer            AS          InPkt    OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.12.1  65400   2471   2473    0    0  20:35:20 Establ
vpn-a.inet.0: 5/5/0
```

Meaning

The sample output for the **show bgp summary instance vpn-a** command shows that the peering session between the CE1 and PE1 routers is established, indicating that the peers are exchanging update messages.

Sample Output

The **show configuration statement-path** command is used to display a specific configuration hierarchy; in this case, the MPLS hierarchy.

```
user@R2-PE1> show configuration protocols mpls
label-switched-path r2-r4 {
    to 192.168.4.1;
    link-protection ;
    primary direct ;
}
path direct {
    10.0.24.2 strict;
}
interface all;
interface fxp0.0 {
    disable;
}
```

Meaning

The sample output for the **show configuration protocols mpls** command shows the current running MPLS configuration on the ingress PE1 router. The configuration include the LSP **r2-r4**, link protection, and the strict primary path **direct**.

Sample Output

The **show mpls lsp** command is used to display summarized information about the configured and active LSPs on a router; in this case, the command shows only the ingress LSPs on the ingress PE1 router.

```
user@R2-PE1> show mpls lsp ingress
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
192.168.4.1 192.168.2.1 Up      0 direct        *      r2-r4
Total 1 displayed, Up 1, Down 0
```

Meaning

The sample output for the **show mpls lsp ingress** command shows that the ingress LSP **r2-r4** is up and following the configured path **direct**.

Sample Output

The **show rsvp session** command is used to display summarized information about active RSVP sessions on a router; in this case, the command shows summarized information about ingress RSVP sessions on the PE1 router

```
user@R2-PE1> show rsvp session ingress
Ingress RSVP: 2 sessions
To          From          State Rt Style Labelin Labelout LSPname
192.168.4.1 192.168.2.1 Up      0 1 SE      -          3 r2-r4
192.168.4.1 192.168.2.1 Up      0 1 SE      - 100064 Bypass->10.0.24.2
Total 2 displayed, Up 2, Down 0
```

Meaning

The sample output for the **show rsvp session ingress** command shows two RSVP sessions are up; the main LSP **r2-r4** and a bypass path protecting the main LSP. Both RSVP sessions are in the Shared Explicit (**SE**) style, creating a shared reservation among for the two paths.

Sample Output

The **show rsvp session ingress detail** command is used to display more detailed information about the two ingress RSVP sessions on the PE1 router.


```
user@R2-PE1> show rsvp session ingress detail
```

```
Ingress RSVP: 2 sessions
```

```
192.168.4.1
```

```
From: 192.168.2.1, LSPstate: Up , ActiveRoute: 0
```

```
LSPname: r2-r4, LSPpath: Primary
```

```
Suggested label received: -, Suggested label sent: -
```

```
Recovery label received: -, Recovery label sent: 3
```

```
Resv style: 1 SE, Label in: -, Label out: 3
```

```
Time left: -, Since: Fri Mar 9 14:05:03 2007
```

```
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
```

```
Port number: sender 1 receiver 63395 protocol 0
```

```
Link protection desired
```

```
Type: Link protected LSP
```

```
PATH rcvfrom: localclient
```

```
Adspec: sent MTU 1500
```

```
Path MTU: received 1500
```

```
PATH sentto: 10.0.24.2 (so-0/0/1.0) 2008 pkts
```

```
RESV rcvfrom: 10.0.24.2 (so-0/0/1.0) 2006 pkts
```

```
Explct route: 10.0.24.2
```

```
Record route: <self> 10.0.24.2
```

```
192.168.4.1
```

```
From: 192.168.2.1, LSPstate: Up, ActiveRoute: 0
```

```
LSPname: Bypass->10.0.24.2
```

```
Suggested label received: -, Suggested label sent: -
```

```
Recovery label received: -, Recovery label sent: 100064
```

```
Resv style: 1 SE, Label in: -, Label out: 100064
```

```
Time left: -, Since: Fri Mar 9 14:05:58 2007
```

```
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
```

```
Port number: sender 1 receiver 63396 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: 10.0.23.14 (fe-0/1/1.0) 2001 pkts
```

```
RESV rcvfrom: 10.0.23.14 (fe-0/1/1.0) 1736 pkts
```

```
Explct route: 10.0.23.14 10.0.34.14
```

```
Record route: <self> 10.0.23.14 10.0.34.14
```

```
Total 2 displayed, Up 2, Down 0
```


Meaning

The sample output for the **show rsvp session ingress detail** command shows the RSVP session for the ingress LSP and the bypass path, which appears as a separate RSVP ingress session for the protected interface **10.0.24.2**. The bypass path is automatically generated. By default, the name appears as **Bypass > interface-address**, where the interface address is the next downstream router's interface (**10.0.24.2**). The explicit route **10.0.23.14 10.0.34.14** for the session shows **R3** as the transit node and **R4** as the egress node.

Sample Output

The **show route table routing-table-name** command is used to display information about a particular routing table. In this case, the **vpn-a.inet.0** routing table.

```
user@R2-PE1> show route table vpn-a 192.168.5.1 detail
vpn-a.inet.0: 9 destinations, 13 routes (9 active, 0 holddown, 0 hidden)
192.168.5.0/24 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
              Route Distinguisher: 192.168.4.1:4
              Next-hop reference count: 11
              Source: 192.168.4.1
              Next hop: via so-0/0/1.0 weight 0x1, selected
              Label-switched-path r2-r4
              Label operation: Push 100160
              Next hop: 10.0.23.14 via fe-0/1/1.0 weight 0x8001
              Label-switched-path r2-r4
              Label operation: Push 100160, Push 100064(top)
              Protocol next hop: 192.168.4.1
              Push 100160
              Indirect next hop: 8791000 262142
              State: <Secondary Active Int Ext>
              Local AS: 65432 Peer AS: 65432
              Age: 1d 5:22:31      Metric2: 1
              Task: BGP_65432.192.168.4.1+2056
              Announcement bits (1): 0-KRT
              AS path: 65400 I
              Communities: target:65432:100
              VPN Label: 100160
              Localpref: 100
              Router ID: 192.168.4.1
              Primary Routing Table bgp.l3vpn.0
```

Meaning

The sample output for the **show route table vpn-a 192.168.5.1 detail** command shows routes associated with the remote PE-CE location as indicated by the loopback address of the PE2 router **192.168.5.1**. In this case, there are different next hops with unequal weights (**0x1** and **0x8001**) associated with the remote location. For correct traffic protection, those two routes must appear in the forwarding table.

Sample Output

The **show route forwarding-table** command displays the route entries in the kernel's forwarding table. This is the version of the forwarding table in the Routing Engine. The Routing Engine copies this table to the Packet Forwarding Engine. In this case, the set of routes installed in the forwarding table to verify that the routing protocol process (rpd) has relayed the correct information to the forwarding table for the specified destination.

```
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1      extensive
Routing table: vpn-a.inet [Index 2]
Internet:

Destination:  192.168.5.0/24
Route type:   user
Route reference: 0                      Route interface-index: 0
Flags: sent to PFE, prefix load balance
Next-hop type: indirect                  Index: 262142   Reference: 2
Next-hop type: Push 100160
Next-hop interface: so-0/0/1.0
```

Meaning

The sample output for the **show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive** command shows only one next hop **so-0/0/1.0** is installed in the forwarding table, indicating that the information in the forwarding table is not correct. We would expect to see the same paths installed in the forwarding table as appear in the routing table in the output for the **show route table vpn-a 192.168.5.1 detail**.

Solution

The solution is to enable load-balancing and ensure that multiple next-hop forwarding table entries appear in the forwarding table for each destination. The forwarding-table entries can be an incoming MPLS label or an IP destination.

A load-balancing policy applied to the forwarding-table is the same mechanism required for ECMP (equal-cost multipath) load-balancing to install multiple next-hops into the forwarding-table. The extra paths installed for local repair are not used for load-balancing, because the paths are differently weighted, as demonstrated in the sample output for the **show routing table** and the **show route forwarding-table** commands.

NOTE: The load-balancing policy must be applied to all provider (P) and provider-edge (PE) routers that are required to support local repair.

The following sample output shows an example load-balancing configuration and the commands used to verify that the required two next-hop entries appear in the forwarding table.

Sample Output

Use the following two **show configuration statement-path** commands to display a specific configuration hierarchy; in this case, policy-options and routing-options.

```
user@R2-PE1> show configuration policy-options
policy-statement lbpf {
    then {
        load-balance per-packet ;
    }
}
[...Output truncated...]

user@R2-PE1> show configuration routing-options
static {
    [...Output truncated...]
    route 100.100.1.0/24 reject;
}
router-id 192.168.2.1;
route-distinguisher-id 192.168.2.1;
autonomous-system 65432;
forwarding-table {
    export lbpf ;
}
```

Meaning

The sample output for the **show configuration policy-options** and **show configuration routing-options** commands shows the two parts required to configure a load balancing policy. The **lbpf** policy includes the **load-balance per-packet** statement. The policy is then applied at the **[edit routing options forwarding-table]** hierarchy level with the **export lbpf** statement. Enabling load balancing results in the export of routes from the routing table to the forwarding table, and a solution to the problem.

NOTE: The **load-balance per-packet** statement is named *per-packet* for historical reasons. When the Packet Forwarding Engine was an IP Processor-1 (before Junos 4.0), Junos supported only per-packet load balancing. When the IP Processor-II was introduced the behavior was changed to per-flow load balancing without changing the statement.

Sample Output

Use the **show route forwarding-table** command to 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. In this case, the option **vpn vpn** is used to display routing table entries for the specified VPN **vpn-a**.

```
user@R2-PE1> show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive
Routing table: vpn-a.inet [Index 2]
Internet:

Destination: 192.168.5.0/24
Route type: user
Route reference: 0                               Route interface-index: 0
Flags: sent to PFE
Next-hop type: indirect                          Index: 262142   Reference: 2
Next-hop type: unilist                           Index: 262146   Reference: 1
Next-hop type: Push 100160
Next-hop interface: so-0/0/1.0   Weight: 0x1
Nexthop: 10.0.23.14
Next-hop type: Push 100160, Push 100064(top)
Next-hop interface: fe-0/1/1.0   Weight: 0x8001
```

Meaning

The sample output for the **show route forwarding-table vpn vpn-a destination 192.168.5.1 extensive** command shows the correct two routes were relayed from the routing table to the forwarding table.

Conclusion

In conclusion, a load balancing policy is required for link protection to work effectively. The principles are the same for the configuration of the **fast reroute** and the **node-link protection** statements.

Router Configurations

The following output shows the configurations of all routers in the network. The **no-more** option entered after the pipe (|) prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output

The following sample output is for the customer edge (CE) 1 router:

```
user@R1-CE1> show configuration | no-more
[...Output truncated...]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.0.12.1/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.143/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.1.1/32;
      }
    }
  }
}
routing-options {
  static {
    /* corporate and alpha net */
    route 172.16.0.0/12 {
      next-hop 192.168.71.254;
      retain;
      no-readvertise;
    }
    /* old lab nets */
    route 192.168.0.0/16 {
      next-hop 192.168.71.254;
      retain;
      no-readvertise;
    }
  }
}
```



```

        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
        route 172.16.0.0/24 reject;
        route 172.16.1.0/24 reject;
        route 172.16.2.0/24 reject;
        route 172.16.3.0/24 reject;
        route 192.168.1.0/24 reject;
    }
    router-id 192.168.1.1;
    autonomous-system 65400;
}
protocols {
    bgp {
        group PE1 {
            type external;
            export stat;
            peer-as 65432;
            neighbor 10.0.12.2;
        }
    }
    ospf {
        traffic-engineering;
        export stat;
        area 0.0.0.0 {
            interface so-0/0/0.0;
            interface lo0.0 {
                passive;
            }
        }
    }
}
policy-options {
    policy-statement stat {
        term 1 {
            from protocol static;
            then accept;
        }
        term 2 {
            then reject;
        }
    }
}

```



```

    }
}

```

Sample Output

The following sample output is for the provider edge (PE) 1 ingress router :

```

user@R2-PE1> show configuration | no-more
[...Output truncated...]
interfaces {
  so-0/0/0 {
    description to-r1;
    unit 0 {
      family inet {
        address 10.0.12.2/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/1 {
    description to-r4;
    unit 0 {
      family inet {
        address 10.0.24.1/30;
      }
      family iso;
      family mpls;
    }
  }
  fe-0/1/1 {
    description to-r3;
    unit 0 {
      family inet {
        address 10.0.23.13/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {

```



```

        address 192.168.70.144/21;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.2.1/32;
        }
    }
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
        route 100.100.1.0/24 reject;
    }
    router-id 192.168.2.1;
    route-distinguisher-id 192.168.2.1;
    autonomous-system 65432;
    forwarding-table {
        export lbpf;
    }
}
protocols {
    rsvp {
        interface fxp0.0 {
            disable;
        }
        interface all {

```



```

        link-protection;
    }
}
mpls {
    label-switched-path r2-r4 {
        to 192.168.4.1;
        link-protection;
        primary direct;
    }
    path via-r3 {
        10.0.23.14 strict;
        10.0.34.14 strict;
    }
    path direct {
        10.0.24.2 strict;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    export send-statics;
    group ibgp {
        type internal;
        local-address 192.168.2.1;
        family inet {
            unicast;
        }
        family inet-vpn {
            unicast;
        }
        export next-hop-self;
        peer-as 65432;
        neighbor 192.168.4.1;
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0 {
            passive;
        }
        interface fe-0/1/1.0;
    }
}

```



```

        interface so-0/0/1.0;
    }
}
policy-options {
    policy-statement lbpf {
        then {
            load-balance per-packet;
        }
    }
    policy-statement next-hop-self {
        from route-type external;
        then {
            next-hop self;
        }
    }
    policy-statement send-statics {
        term statics {
            from {
                route-filter 100.100.1.0/24 exact;
            }
            then accept;
        }
    }
    policy-statement vpna-export {
        term 1 {
            from protocol static;
            then {
                community add vpna-target;
                community add vpna-origin;
                accept;
            }
        }
        term 2 {
            then reject;
        }
    }
    policy-statement vpna-import {
        term 1 {
            from {
                protocol bgp;
                community vpna-target;
            }
            then accept;
        }
    }
}

```



```

    }
    term 2 {
        then reject;
    }
}
community vpna-origin members origin:192.168.2.1:1;
community vpna-target members target:65432:100;
}
routing-instances {
    vpn-a {
        instance-type vrf;
        interface so-0/0/0.0;
        vrf-target {
            import target:65432:100;
            export target:65432:100;
        }
        protocols {
            bgp {
                group CE1 {
                    type external;
                    peer-as 65400;
                    neighbor 10.0.12.1;
                }
            }
        }
    }
}
}

```

Sample Output

The following sample output is for the provider (P) transit router:

```

user@R3-P> show configuration | no-more
[...Output truncated...]
interfaces {
    fe-1/3/0 {
        description to-r4;
        unit 0 {
            family inet {
                address 10.0.34.13/30;
            }
            family iso;
            family mpls;
        }
    }
}

```



```

    }
}
fe-1/3/1 {
    description to-r2;
    unit 0 {
        family inet {
            address 10.0.23.14/30;
        }
        family iso;
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.145/21;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.3.1/32;
        }
        family iso {
            address 49.0004.1921.6800.3001.00;
        }
    }
}
}
routing-options {
    static {
        /* corperate and alpha net */
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        /* old lab nets */
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
    }
}

```



```

        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
    }
    router-id 192.168.3.1;
    autonomous-system 65432;
}
protocols {
    rsvp {
        interface all {
            link-protection;
        }
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        icmp-tunneling;
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface fxp0.0 {
                disable;
            }
            interface all;
        }
    }
}
}

```

Sample Output

The following sample output is for the provider edge (PE) 2 ingress router :


```

user@R4-PE2> show configuration | no-more
[...Output truncated...]
interfaces {
  so-0/0/1 {
    description to-R2;
    unit 0 {
      family inet {
        address 10.0.24.2/30;
      }
      family iso;
      family mpls;
    }
  }
  so-0/0/2 {
    description to-R5-CE2;
    unit 0 {
      family inet {
        address 10.0.45.1/30;
      }
      family iso;
      family mpls;
    }
  }
  fe-0/1/3 {
    description to-R3-P;
    unit 0 {
      family inet {
        address 10.0.34.14/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.146/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 192.168.4.1/32;
      }
    }
  }
}

```



```

        }
    }
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
        route 100.100.4.0/24 reject;
    }
    router-id 192.168.4.1;
    route-distinguisher-id 192.168.4.1;
    autonomous-system 65432;
    forwarding-table {
        export lbpf;
    }
}
protocols {
    rsvp {
        interface fxp0.0 {
            disable;
        }
        interface all {
            link-protection;
        }
    }
    mpls {
        label-switched-path r4-r2 {
            to 192.168.2.1;
        }
        interface all;
    }
}

```



```

        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        export send-statics;
        group ibgp {
            type internal;
            local-address 192.168.4.1;
            family inet {
                unicast;
            }
            family inet-vpn {
                unicast;
            }
            export next-hop-self;
            peer-as 65432;
            neighbor 192.168.2.1;
        }
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface lo0.0 {
                passive;
            }
            interface fe-0/1/3.0;
            interface so-0/0/1.0;
        }
    }
}
policy-options {
    policy-statement lbpf {
        then {
            load-balance per-packet;
        }
    }
    policy-statement next-hop-self {
        from route-type external;
        then {
            next-hop self;
        }
    }
    policy-statement send-statics {

```



```

    term statics {
        from {
            route-filter 100.100.4.0/24 exact;
        }
        then accept;
    }
}
policy-statement vpnb-export {
    term 1 {
        from protocol static;
        then {
            community add vpnb-target;
            community add vpnb-origin;
            accept;
        }
    }
    term 2 {
        then reject;
    }
}
policy-statement vpnb-import {
    term 1 {
        from {
            protocol bgp;
            community vpnb-target;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
community vpnb-origin members origin:192.168.5.1:1;
community vpnb-target members target:65432:100;
}
routing-instances {
    vpn-b {
        instance-type vrf;
        interface so-0/0/2.0;
        vrf-target {
            import target:65432:100;
            export target:65432:100;
        }
        protocols {

```



```

        bgp {
            group CE2 {
                type external;
                peer-as 65400;
                neighbor 10.0.45.2;
            }
        }
    }
}

```

Sample Output

The following sample output is for the customer edge (CE) 2 router:

```

user@R5-CE2> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.0.45.2/30;
            }
        }
    }
    fxp0 {
        unit 0 {
            family inet {
                address 192.168.70.147/21;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 192.168.5.1/32;
            }
            family iso {
                address 49.0004.1921.6800.5001.00;
            }
        }
    }
}

```



```

routing-options {
  graceful-restart;
  static {
    /* corporate and alpha net */
    route 172.16.0.0/12 {
      next-hop 192.168.71.254;
      retain;
      no-readvertise;
    }
    /* old lab nets */
    route 192.168.0.0/16 {
      next-hop 192.168.71.254;
      retain;
      no-readvertise;
    }
    route 0.0.0.0/0 {
      discard;
      retain;
      no-readvertise;
    }
    route 172.16.0.0/24 reject;
    route 172.16.1.0/24 reject;
    route 172.16.2.0/24 reject;
    route 172.16.3.0/24 reject;
    route 192.168.5.0/24 reject;
  }
  router-id 192.168.5.1;
  autonomous-system 65400;
}

protocols {
  bgp {
    group PE2 {
      type external;
      export stat;
      peer-as 65432;
      neighbor 10.0.45.1;
    }
  }
  ospf {
    traffic-engineering;
    export stat;
    area 0.0.0.0 {
      interface so-0/0/2.0;
      interface lo0.0 {

```



```

        passive;
    }
}
}
}
policy-options {
    policy-statement stat {
        term 1 {
            from protocol static;
            then accept;
        }
        term 2 {
            then reject;
        }
    }
}
}
```

Problem Establishing a GRE Tunnel Checklist

Problem

Description: This checklist provides the links and commands for troubleshooting a case study about a problem establishing a Generalized Multiprotocol Label Switching (GMPLS) label-switched path (LSP). Specifically, the configuration of the data channel is incorrect because the configuration includes different interface types at both ends of the tunnel. The principles and solution used in this case study also apply to control channel configuration.

The checklist includes the links to a brief summary of GRE tunnels within the context of GMPLS, an example network scenario, and more detailed information about the commands used to troubleshoot and resolve the problem.

The troubleshooting process described in this case study should not be followed rigidly; it is a basis from which you can develop your own process to suit your particular situation. (See [Table 64 on page 2257](#))

Solution

Table 64: Problem Establishing a GRE Tunnel Checklist

Tasks	Command or Action
“Troubleshooting GMPLS and GRE Tunnel” on page 2258	

Table 64: Problem Establishing a GRE Tunnel Checklist (continued)

Tasks	Command or Action
"Symptom" on page 2261	<pre>show mpls lsp show rsvp session</pre>
"Cause" on page 2262	The cause of the problem with the GMPLS LSP is the configuration of different interface types at both ends of the GMPLS data channel.
"Troubleshooting Commands" on page 2262	<pre>show mpls lsp extensive show rsvp session detail show link-management peer show link-management te-link show configuration protocols mpls monitor start filename show log filename</pre>
"Solution" on page 2269	<p>Configure both ends of the data channel with the same switching type.</p> <pre>show configuration protocols link-management show mpls lsp show link-management te-link</pre>
"Conclusion" on page 2271	Both ends of a GMPLS data must be the same encapsulation or interface type.
"Router Configurations" on page 2271	<code>show configuration no-more</code>

Troubleshooting GMPLS and GRE Tunnel

Problem

Description: The logical control channel for GMPLS must be a point-to-point link and must have some form of IP reachability. On broadcast interfaces or when there are multiple hops between control channel peers, use a GRE tunnel for the control channel. For more detailed information on GMPLS and GRE tunnels see the *Junos MPLS Applications Configuration Guide* and the *Junos User Guide*.

A tunnel PIC is *not* required to configure a GRE tunnel for the GMPLS control channel. Instead, use the software-based **gre** interface, rather than the hardware-based **gr-fpc/pic/port** interface.



CAUTION: Due to restrictions to the software-based **gre** interface, the GMPLS control channel is the only supported use of the software-based **gre** interface. Any other use is expressly unsupported and might cause an application failure.

The following example shows a basic **gre** interface configuration. In this case, the tunnel source is the loopback address of the local router and the destination address is the loopback destination of the remote router. Traffic that has a next hop of the tunnel destination will use the tunnel. The tunnel is not automatically used by all the traffic passing through the interface. Only traffic with the tunnel destination as the next hop uses the tunnel.

Sample Output

```
user@R1> show configuration interfaces
[...Output truncated...]
gre {
    unit 0 {
        tunnel {
            source 10.0.12.13;
            destination 10.0.12.14;
        }
        family inet {
            address 10.35.1.6/30;
        }
        family mpls;
    }
}
```

Sample Output

The following sample output for the **show interfaces** command shows the encapsulation type and header, the maximum speed, packets through the logical interface, the destination, and logical address.

```
user@R1> show interfaces gre
Physical interface: gre, Enabled, Physical link is Up
Interface index: 10, SNMP ifIndex: 8
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: Unlimited
Device flags      : Present Running
Interface flags: Point-To-Point SNMP-Traps
Input packets    : 0
Output packets   : 0
```



```

Logical interface gre.0 (Index 70) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 0x4000
    IP-Header 10.0.12.14:10.0.12.13:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
    Input packets : 171734
    Output packets: 194560
    Protocol inet, MTU: 1476
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.35.1.4/30, Local: 10.35.1.6, Broadcast: 10.35.1.7
    Protocol mpls, MTU: 1464
      Flags: None


```

The following are various requirements when you configure a GMPLS LSP using a GRE tunnel:

- The data channel must start and end on the same type of interface.
- The control channel can be a GRE tunnel that starts and ends on the same or different interface type.
- The GRE tunnel must be configured indirectly with the **peer-interface *peer-name*** statement at the **[edit protocols ospf]** hierarchy level.
- The GRE interface must be disabled at the **[edit protocols ospf]** and **[edit protocols rsvp]** hierarchy levels.
- Data and control channels must be defined correctly in the LMP configuration .
- It is optional to disable Constrained Shortest Path First (CSPF) with the **no-cspf** statement.

This case focuses on the incorrect configuration of the endpoints of the GRE tunnel. However, you can use a similar process and commands to diagnose other GRE tunnel problems. [Figure 169 on page 2260](#) illustrates a network topology with MPLS tunneled through a GRE interface.

Figure 169: GMPLS Network Topology



The MPLS network topology in [Figure 169 on page 2260](#) shows Juniper Networks routers configured with a GRE tunnel that consists of the following components:

- A strict GMPLS LSP path from the ingress router to the egress router.
- On the ingress router, CSPF disabled with the **no-cspf** statement at the [edit protocol mpls label-switched-path *lsp-name*] hierarchy level.
- Traffic-engineering links and control channels within the **peer** statement at the [edit protocols link-management] hierarchy level on all routers.
- OSPF and OSPF traffic engineering configured on all routers.
- A reference to the **peer-interface** in both OSPF and RSVP on all routers.
- A switching-type problem between **R2** and **R3**.

Symptom

The LSP in the network shown in [Figure 169 on page 2260](#) is down, as indicated by the output from the **show mpls lsp** and **show rsvp session** commands, which display very similar information. The **show mpls lsp** command shows all LSPs configured on the router, as well as all transit and egress LSPs. The **show rsvp session** command displays summary information about RSVP sessions. You can use either command to verify the state of the LSP. In this case, LSP **gmpls-r1-to-r3** is down (**Dn**).

Sample Output

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
192.168.4.1 192.168.1.1 Dn  0 -      gmpls-r1-to-r3
Bidir
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R1> show rsvp session
Ingress RSVP: 1 sessions
To          From          State  Rt Style Labelin Labelout LSPname
192.168.4.1 192.168.1.1 Dn    0 0 -    -    - gmpls-r1-to-r3
Bidir
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
```

Cause

The cause of the problem with the GMPLS LSP is the configuration of different interface types at both ends of the GMPLS data channel.

Troubleshooting Commands

The Junos OS includes commands that are useful when troubleshooting a problem. This topic provides a brief description of each command, followed by sample output, and a discussion of the output in relation to the problem.

You can use the following commands when troubleshooting a GMPLS problem:

```
user@host> show mpls lsp extensive
user@host> show rsvp session detail
user@host> show link-management peer
user@host> show link-management te-link
user@host> show configuration protocols mpls
user@host> monitor start filename
user@host> show log filename
```

Sample Output

Use the show mpls lsp extensive command on transit router R1 to display detailed information about all LSPs transiting, terminating, and configured on the router.

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.4.1
  From: 192.168.1.1, State: Dn, ActiveRoute: 0, LSPname: gmpls-r1-to-r3
  Bidirectional
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: SDH/SONET, Switching type: PSC-1, GPID: IPv4
  Primary   pl                               State: Dn
    SmartOptimizeTimer: 180
8 Dec 20 18:08:02 192.168.4.1: MPLS label allocation failure [3 times]
  7 Dec 20 18:07:53 Originate Call
  6 Dec 20 18:07:53 Clear Call
  5 Dec 20 18:07:53 Deselected as active
```



```

    4 Dec 20 18:06:13 Selected as active path
    3 Dec 20 18:06:13 Record Route:  100.100.100.100 93.93.93.93
    2 Dec 20 18:06:13 Up
    1 Dec 20 18:06:13 Originate Call
Created: Wed Dec 20 18:06:12 2006
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output for the **show mpls lsp extensive** command shows an error message (**MPLS label allocation failure**) in the log section of the output. This LSP event indicates that the MPLS protocol or the **family mpls** statement were not configured properly. When the LSP event is preceded by an IP address, the address is typically the router that has the MPLS configuration error. In this case, it appears that the router with the **lo0** address of **192.168.4.1 (R3)** has an MPLS configuration error.

Sample Output

Use the **show rsvp session detail** command to display detailed information about RSVP sessions.

```

user@R1> 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: Wed Dec 20 18:07:53 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 (tester2) 3 pkts
    Explt route: 100.100.100.100 93.93.93.93

```


Record route: <self> ...incomplete

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

Meaning

The sample output for the **show rsvp session detail** command shows that LSP **gmpls-r1-to-r3** is down (**LSPstate: Dn**). The route record is incomplete, indicating a problem with the explicit route **100.100.100.100 93.93.93.93**. The address **100.100.100.100** is the data channel on **R2 so-0/0/0**, and the address **93.93.93.93** is the data channel on **R3**.

Sample Output

Use the **show link-management peer** command to display MPLS peer link information.

```
user@R1> show link-management peer
```

Peer name: tester2, System identifier: 48428

State: Up, Control address: 10.35.1.5

Control-channel	State
gre.0	Active

TE links:

tester2

```
user@R2> show link-management peer
```

Peer name: tester2, System identifier: 48428

State: Up, Control address: 10.35.1.6

Control-channel	State
gre.0	Active

TE links:

te-tester2

Peer name: tester3 , System identifier: 48429

State: Up , Control address: 10.35.1.2

Control-channel	State
gre.1	Active

TE links:

te-tester3


```

user@R3> show link-management peer
Peer name: tester3, System identifier: 48429
State: Up, Control address: 10.35.1.1
Control-channel          State
gre.0                    Active
TE links:
te-tester3

```

Meaning

The sample output from all routers in the example network in [Figure 169 on page 2260](#) for the **show link-management peer** command shows that all control channels are up and active. A detailed analysis of the output shows the following information:

- Name of the peer, **tester2** or **tester3**, which is the same on neighboring routers for ease of troubleshooting.
- Internal identifier for the peer, **48428** for **tester2** and **48429** for **tester3**. The internal identifier is a range of values from 0 through 64,000.
- The state of the peer, which can be up or down. In this case, all peers are up.
- The address to which a control channel is established, for example, **10.35.1.5**.
- The state of the control channel, which can be up, down, or active.
- The traffic-engineered links that are managed by their peer, indicating that control channel **gre.0** is managed by **tester3**.

Sample Output

Use the **show link-management te-link** command to display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

```

user@R1> show link-management te-link
TE link name:  tester2, State: Up
Local identifier: 2005, Remote identifier: 21253, Local address: 90.90.90.90,
Remote address: 100.100.100.100,
Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
Available bandwidth: 0bps

```

Name	State	Local ID	Remote ID	Bandwidth	Used	LSP-name
so-0/0/0	Up	21253	21253	155.52Mbps	Yes	gmpls-r1-to-r3

```

user@R2> show link-management te-link

```



```

TE link name:  te-tester2, State: Up
  Local identifier: 7002, Remote identifier: 22292, Local address: 100.100.100.100,
  Remote address: 90.90.90.90,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
  bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name          State Local ID  Remote ID      Bandwidth  Used   LSP-name
    so-0/0/0      Up          21253         21253        155.52Mbps  Yes   gmpls-r1-to-r3
TE link name:  te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254, Local address: 103.103.103.103,
  Remote address: 93.93.93.93,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
  bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
  Available bandwidth: 0bps
    Name          State Local ID  Remote ID      Bandwidth  Used   LSP-name
    so-0/0/1      Up          21252         21252        155.52Mbps  Yes   gmpls-r1-to-r3

user@R3> show link-management te-link
TE link name:  te-tester3, State: Up
  Local identifier: 7003, Remote identifier: 21254,  Local address: 93.93.93.93,
  Remote address: 103.103.103.103,
  Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 0bps, Maximum bandwidth:
  0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name          State Local ID  Remote ID      Bandwidth Used   LSP-name
    so-0/0/1      Dn          21252         21252        155.52Mbps  No

```

Meaning

The sample output for the **show link-management te-link** command issued on the three routers in the network in [Figure 169 on page 2260](#) shows the resources allocated to the traffic-engineered links **te-tester2** and **te-tester3**. The resources are the SONET interfaces **so-0/0/0** and **so-0/0/1**. On **R1** and **R2**, the SONET interfaces are used for the LSP **gmpls-r1-to-r3**, as indicated by **Yes** in the **Used** field. However, the SONET interface **so-0/0/1** on **R3** is down (**Dn**) and is not used for the LSP (**Used No**). Further investigation is required to discover why the SONET interface on **R3** is down.

Sample Outut

Use the `show log filename` command to display the contents of the specified log file. In this case, the log file `rsvp.log` is configured at the `[edit protocols rsvp traceoptions]` hierarchy level. When the log file is configured, you must issue the `monitor start filename` command to begin logging messages to the file.

```
user@R1> show configuration protocols rsvp
traceoptions {
    file rsvp.log size 3m world-readable;
    flag state detail;
    flag error detail;
    flag packets detail;
}

user@R1> monitor start rsvp.log
```

NOTE: The **find Error** option entered after the pipe (|) searches the output for an instance of the term *Error*.

Sample Output

```
user@R3>
show log rsvp.log | find Error

Dec 28 17:23:32 Error Len 20 Session preempted flag 0 by 192.168.4.1 TE-link
103.103.103.103
[...Output truncated...]
Dec 28 17:23:32 RSVP new resv state,session 192.168.4.1(port/tunnel ID 46115 Ext-ID
192.168.1.1)Proto 0
Dec 28 17:23:32          RSVP-LMP reset LMP request for gmpls-r1-to-r3
Dec 28 17:23:32          RSVP->LMP request - resource for LSP gmpls-r1-to-r3
Dec 28 17:23:32      LMP->RSVP resource request gmpls-r1-to-r3 failed cannot find resource encoding type
SDH/SONET remote label 21252 bandwidth bw[0
Dec 28 17:23:32      RSVP-LMP reset LMP request for gmpls-r1-to-r3
Dec 28 17:23:32 RSVP originate PathErr 192.168.4.1->192.168.2.1 MPLS label allocation failure LSP
gmpls-r1-to-r3(2/46115)
Dec 28 17:23:32 RSVP send PathErr 192.168.4.1->192.168.2.1 Len=196 tester3
Dec 28 17:23:32 Session7 Len 16 192.168.4.1(port/tunnel ID 46115 Ext-ID
192.168.1.1) Proto 0
Dec 28 17:23:32 Hop          Len 20 192.168.4.1/0x086e4770 TE-link 103.103.103.103
Dec 28 17:23:32 Error       Len 20 MPLS label allocation failure flag 0 by
```



```

192.168.4.1 TE-link 103.103.103.103
Dec 28 17:23:32   Sender7   Len 12 192.168.1.1(port/lsp ID  2)
Dec 28 17:23:32   Tspec      Len 36 rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
Dec 28 17:23:32   ADspec     Len 48 MTU 1500
Dec 28 17:23:32   RecRoute   Len 20  103.103.103.103 90.90.90.90
Dec 28 17:23:32   SuggLabel  Len  8  21252
Dec 28 17:23:32   UpstrLabel Len  8  21252

```

Meaning

The sample output from the egress router **R3** for the **show log rsvp.log** command is a snippet taken from the log file. The snippet shows a Link Management Protocol (LMP) resource request for the LSP **gmpls-r1-to-r3**. The request has problems with the encoding type (SDH/SONET), indicating a possible error with the SONET interface connecting **R2** and **R3**. Further investigation of the configuration of the LMP on **R2** and **R3** is required.

Sample Output

Use the show configuration *statement-path* command to display a specific configuration hierarchy; in this instance, link-management.

```

user@R2> show configuration protocols link-management
te-link te-tester2 {
    local-address 100.100.100.100;
    remote-address 90.90.90.90;
    remote-id 22292;
    interface so-0/0/0 {
        local-address 100.100.100.100;
        remote-address 90.90.90.90;
        remote-id 21253;
    }
}
te-link te-tester3 {
    local-address 103.103.103.103;
    remote-address 93.93.93.93;
    remote-id 21254;
    interface so-0/0/1 {
        local-address 103.103.103.103;
        remote-address 93.93.93.93;
        remote-id 21252;
    }
}
peer tester2 {
    address 10.35.1.6;
}

```



```

        control-channel gre.0;
        te-link te-tester2;
    }
    peer tester3 {
        address 10.35.1.2;
        control-channel gre.1;
        te-link te-tester3;
    }

user@R3> show configuration protocols link-management
te-link te-tester3 {
    local-address 93.93.93.93;
    remote-address 103.103.103.103;
    remote-id 21254;
}
    interface at-0/3/1 {
        local-address 93.93.93.93;
        remote-address 103.103.103.103;
        remote-id 21252;
    }
}
peer tester3 {
    address 10.35.1.1;
    control-channel gre.0;
    te-link te-tester3;
}

```

Meaning

The sample output from transit router **R2** and ingress router **R3** for the **show configuration protocols link-management** command shows that the interface type on the two routers is different. The resource allocated to **te-tester3** on transit router **R2** is a SONET interface, while the resource allocated to **te-tester3** on egress router **R3** is an ATM interface. The interface type on each end of the data or control channels must be of the same type. In this case, both ends should be SONET or ATM.

Solution

Solution

The solution to the problem of different interface or encapsulation types at either end of the GMPLS LSP is to make sure that the interface type is the same at both ends. In this case, the ATM interface was deleted from the link-management configuration on **R3**, and a SONET interface was configured instead.

The following commands illustrate the correct configuration and commands to verify that the GMPLS LSP is up and using the data channel:


```

user@R3> show configuration protocols link-management
user@R3> show mpls lsp
user@R3> show link-management te-link

```

Sample Output

```

user@R3> show configuration protocols link-management
te-link te-tester3 {
    local-address 93.93.93.93;
    remote-address 103.103.103.103;
    remote-id 21254;
    interface so-0/0/1 { # SONET interface replaces the incorrect ATM interface
        local-address 93.93.93.93;
        remote-address 103.103.103.103;
        remote-id 21252;
    }
}
peer tester3 {
    address 10.35.1.1;
    control-channel gre.0;
    te-link te-tester3;
}

user@R3> show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Egress LSP: 1 sessions

```

To	From	State	Rt	Style	Labelin	Labelout	LSPname
192.168.4.1	192.168.1.1	Up	0	1 FF	21252	-	gmpls-r1-to-r3

```

Bidir
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show link-management te-link
TE link name: te-tester3, State: Up
Local identifier: 7003, Remote identifier: 21254, Local address: 93.93.93.93,
Remote address: 103.103.103.103,
Encoding: SDH/SONET, Switching: PSC-1, Minimum bandwidth: 155.52Mbps, Maximum
bandwidth: 155.52Mbps, Total bandwidth: 155.52Mbps,
Available bandwidth: 0bps

```


Name	State	Local ID	ID	Remote ID	Bandwidth Used	LSP-name
so-0/0/1	Up	21252	21252	155.52Mbps	Yes	gmpls-r1-to-r3

Meaning

The sample output for the **show protocols link-management**, **show mpls lsp**, and **show link-management te-link** commands from ingress router **R3** show that the problem is solved. LMP is correctly configured, and the LSP **gmpls-r1-to-r3** is up and using the data channel **so-0/0/1**.

Conclusion

In conclusion, both ends of a GMPLS data channel must be the same encapsulation or interface type. This case illustrates the correct configuration of the data channel. The principles are the same for the control channel.

Router Configurations

Output that shows the configurations of the ingress router in the network. The **no-more** option entered after the pipe (|) prevents the output from being paginated if the output is longer than the length of the terminal screen.

Sample Output

The following sample output is for ingress router R1:

```
user@R1> show configuration | no-more
[...Output truncated...]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        address 10.0.12.1/32 {
          destination 10.0.12.2;
        }
      }
      family mpls;
    }
  }
  fe-0/1/0 {
    unit 0 {
      family inet {
        address 10.0.12.13/30;
      }
      family mpls;
    }
  }
}
```



```

fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.143/21;
        }
    }
}
gre {
    unit 0 {
        tunnel {
            source 10.0.12.13;
            destination 10.0.12.14;
        }
        family inet {
            address 10.35.1.6/30;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.1.1/32;
        }
    }
}
}
routing-options {
    static {
        /* corporate and alpha net */
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        /* old lab nets */
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
        }
    }
}

```



```

        no-readvertise;
    }
}
router-id 192.168.1.1;
autonomous-system 65432;
}
protocols {
    rsvp {
        traceoptions {
            file rsvp.log size 3m world-readable;
            flag state detail;
            flag error detail;
            flag packets detail;
        }
        interface fxp0.0 {
            disable;
        }
        interface all;
        interface lo0.0;
        interface gre.0 {
            disable;
        }
        peer-interface tester2;
    }
    mpls {
        label-switched-path gmpls-r1-to-r3 {
            from 192.168.1.1;
            to 192.168.4.1;
            lsp-attributes {
                switching-type psc-1;
                encoding-type sonet-sdh;
            }
            no-cspf;
            primary p1;
        }
        path p1 {
            100.100.100.100 strict;
            93.93.93.93 strict;
        }
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {

```



```

        interface lo0.0;
        interface fe-0/1/0.0;
        interface fxp0.0 {
            disable;
        }
        interface gre.0 {
            disable;
        }
        peer-interface tester2;
    }
}
link-management {
    te-link tester2 {
        local-address 90.90.90.90;
        remote-address 100.100.100.100;
        remote-id 21253;
        interface so-0/0/0 {
            local-address 90.90.90.90;
            remote-address 100.100.100.100;
            remote-id 21253;
        }
    }
    peer tester2 {
        address 10.35.1.5;
        control-channel gre.0;
        te-link tester2;
    }
}
}

```

Sample Output

The following sample output is for transit router R2:

```

user@R2>show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/0 {
        unit 0 {
            family inet {
                address 10.0.12.2/32 {
                    destination 10.0.12.1;
                }
            }
        }
    }
}

```



```

        }
        family mpls;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.0.24.1/32 {
                destination 10.0.24.2;
            }
        }
        family mpls;
    }
}
fe-0/1/0 {
    unit 0 {
        family inet {
            address 10.0.12.14/30;
        }
        family mpls;
    }
}
fe-0/1/2 {
    unit 0 {
        family inet {
            address 10.0.24.13/30;
        }
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.144/21;
        }
    }
}
gre {
    unit 0 {
        tunnel {
            source 10.0.12.14;
            destination 10.0.12.13;
        }
        family inet {

```



```

        address 10.35.1.5/30;
    }
    family mpls;
}
unit 1 {
    tunnel {
        source 10.0.24.13;
        destination 10.0.24.14;
    }
    family inet {
        address 10.35.1.1/30;
    }
    family mpls;
}
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.2.1/32;
        }
    }
}
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 0.0.0.0/0 {
            discard;
            retain;
            no-readvertise;
        }
    }
}
router-id 192.168.2.1;
autonomous-system 65432;
}

```



```

protocols {
  rsvp {
    traceoptions {
      file rsvp.log size 3m world-readable;
      flag packets detail;
      flag state detail;
      flag error detail;
    }
    interface fxp0.0;
    interface lo0.0;
    interface all;
    interface gre.0 {
      disable;
    }
    peer-interface tester2;
    peer-interface tester3;
  }
  mpls {
    interface all;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface lo0.0;
      interface fxp0.0 {
        disable;
      }
      interface gre.0 {
        disable;
      }
      interface fe-0/1/0.0;
      interface fe-0/1/2.0;
      interface gre.1 {
        disable;
      }
      peer-interface tester2;
      peer-interface tester3;
    }
  }
  link-management {
    te-link te-tester2 {
      local-address 100.100.100.100;
      remote-address 90.90.90.90;
      remote-id 22292;
    }
  }
}

```



```

        interface so-0/0/0 {
            local-address 100.100.100.100;
            remote-address 90.90.90.90;
            remote-id 21253;
        }
    }
    te-link te-tester3 {
        local-address 103.103.103.103;
        remote-address 93.93.93.93;
        remote-id 21254;
        interface so-0/0/1 {
            local-address 103.103.103.103;
            remote-address 93.93.93.93;
            remote-id 21252;
        }
    }
    peer tester2 {
        address 10.35.1.6;
        control-channel gre.0;
        te-link te-tester2;
    }
    peer tester3 {
        address 10.35.1.2;
        control-channel gre.1;
        te-link te-tester3;
    }
}

```

Sample Output

The following sample output is for egress router R3:

```

user@R3> show configuration | no-more
[...Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.0.24.2/32;
            }
            family mpls;
        }
    }
}

```



```

}
fe-0/1/2 {
    unit 0 {
        family inet {
            address 10.0.24.14/30;
        }
        family mpls;
    }
}
fxp0 {
    unit 0 {
        family inet {
            address 192.168.70.146/21;
        }
    }
}
gre {
    unit 0 {
        tunnel {
            source 10.0.24.14;
            destination 10.0.24.13;
        }
        family inet {
            address 10.35.1.2/30;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.4.1/32;
        }
    }
}
}
routing-options {
    static {
        route 172.16.0.0/12 {
            next-hop 192.168.71.254;
            retain;
            no-readvertise;
        }
        route 192.168.0.0/16 {

```



```

        next-hop 192.168.71.254;
        retain;
        no-readvertise;
    }
    route 0.0.0.0/0 {
        discard;
        retain;
        no-readvertise;
    }
}
router-id 192.168.4.1;
autonomous-system 65432;
}
protocols {
    rsvp {
        traceoptions {
            file rsvp.log size 3m world-readable;
            flag packets detail;
            flag error;
            flag state;
            flag lmp;
        }
        interface fxp0.0 {
            disable;
        }
        interface all;
        interface lo0.0;
        interface gre.0 {
            disable;
        }
        peer-interface tester3;
    }
    mpls {
        interface all;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface fxp0.0 {
                disable;
            }
            interface fe-0/1/2.0;
            interface gre.0 {
                disable;
            }
        }
    }
}

```



```

        }
        interface lo0.0;
        peer-interface tester3;
    }
}
link-management {
    te-link te-tester3 {
        local-address 93.93.93.93;
        remote-address 103.103.103.103;
        remote-id 21254;
        interface so-0/0/1 {
            local-address 93.93.93.93;
            remote-address 103.103.103.103;
            remote-id 21252;
        }
    }
    peer tester3 {
        address 10.35.1.1;
        control-channel gre.0;
        te-link te-tester3;
    }
}
}

```

Verify Protocol Families

Purpose

If a logical interface does not have MPLS enabled, it cannot perform MPLS switching. This step allows you to quickly determine which interfaces are configured with MPLS and other protocol families.

Action

To verify the protocol families configured on the routers in your network, enter the following Junos OS CLI operational mode command:

```
user@host> show interfaces terse
```


Sample Output 1

user@R1> show interfaces terse

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.12.1/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.15.1/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.13.1/30	
			iso		
			mpls		
so-0/0/3	up	down			

user@R2> show interfaces terse

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.12.2/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.23.1/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.26.1/30	
			iso		
			mpls		
so-0/0/3	up	up			
so-0/0/3.0	up	up	inet	10.1.24.1/30	
			iso		
			mpls		

user@R3> show interfaces terse

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.34.1/30	
			iso		
			mpls		


```

so-0/0/1          up    up
so-0/0/1.0        up    up    inet  10.1.23.2/30
                    iso
                    mpls

so-0/0/2          up    up
so-0/0/2.0        up    up    inet  10.1.13.2/30
                    iso
                    mpls

so-0/0/3          up    up
so-0/0/3.0        up    up    inet  10.1.36.1/30
                    iso
                    mpls

```

user@R4> **show interfaces terse**

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.34.2/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.46.1/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.45.1/30	
			iso		
			mpls		
so-0/0/3	up	up			
so-0/0/3.0	up	up	inet	10.1.24.2/30	
			iso		
			mpls		

user@R5> **show interfaces terse**

Interface	Admin	Link	Proto	Local	Remote
so-0/0/0	up	up			
so-0/0/0.0	up	up	inet	10.1.56.1/30	
			iso		
			mpls		
so-0/0/1	up	up			
so-0/0/1.0	up	up	inet	10.1.15.2/30	
			iso		
			mpls		
so-0/0/2	up	up			
so-0/0/2.0	up	up	inet	10.1.45.2/30	


```

iso
mpls
so-0/0/3          up    down

user@R6> show interfaces terse
Interface          Admin Link Proto Local          Remote
so-0/0/0           up    up
so-0/0/0.0         up    up   inet  10.1.56.2/30
                  iso
                  mpls
so-0/0/1           up    up
so-0/0/1.0         up    up   inet  10.1.46.2/30
                  iso
                  mpls
so-0/0/2           up    up
so-0/0/2.0         up    up   inet  10.1.26.2/30
                  iso
                  mpls
so-0/0/3           up    up
so-0/0/3.0         up    up   inet  10.1.36.2/30
                  iso
                  mpls

```

Sample Output 2

```

user@R6> show interfaces terse
Interface          Admin Link Proto Local          Remote
so-0/0/0           up    up
so-0/0/0.0         up    up   inet  10.1.56.2/30
                  iso
                  mpls
so-0/0/1           up    up
so-0/0/1.0         up    up   inet  10.1.46.2/30
                  iso
                  mpls
so-0/0/2           up    up
so-0/0/2.0         up up inet 10.1.26.2/30
                  iso #The mpls statement is missing.
so-0/0/3           up    up
so-0/0/3.0         up    up   inet  10.1.36.2/30

```



```
iso
mpls
```

Meaning

Sample Output 1 shows the interface, the administrative status of the link (**Admin**), the data link layer status of the link (**Link**), the protocol families configured on the interface (**Proto**), and the local and remote addresses on the interface.

All interfaces on all routes in the network shown in [Figure 167 on page 2218](#) are administratively enabled and functioning at the data link layer with MPLS and IS-IS, and have an **inet** address. All are configured with an IPv4 protocol family (**inet**), and have the IS-IS (**iso**) and MPLS (**mpls**) protocol families configured at the `[edit interfaces type-fpc/pic/port unit number]` hierarchy level.

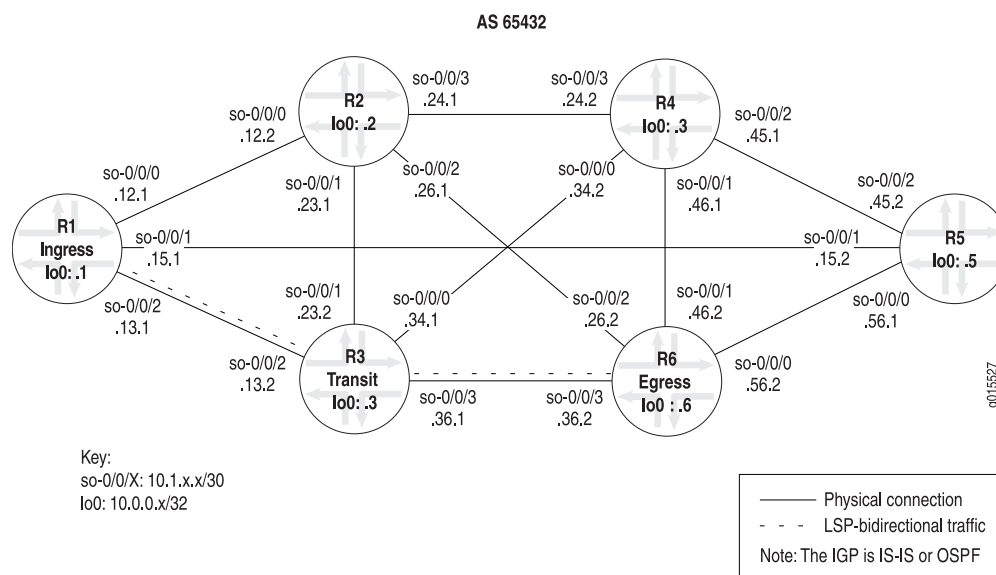
Sample Output 2 shows that interface `so-0/0/2.0` on **R6** does not have the **mpls** statement included at the `[edit interfaces type-fpc/pic/port unit number]` hierarchy level.

Determining LSP Status

Display detailed information about Resource Reservation Protocol (RSVP) objects and the label-switched path (LSP) history to pinpoint a problem with the LSP.

[Figure 170 on page 2285](#) illustrates the network topology used in this topic.

Figure 170: MPLS Network Topology



To determine the LSP state, follow these steps:

1. [Check the Status of the LSP | 2286](#)
2. [Display Extensive Status About the LSP | 2287](#)

Check the Status of the LSP

Purpose

Display the status of the label-switched pathe (LSP).

Action

To determine the LSP status, on the ingress router, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show mpls lsp
```

Sample Output

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Up    1      *   R1-to-R6
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt  Style  Labelin Labelout LSPname
10.0.0.1    10.0.0.6    Up    0 1 FF   3    - R6-to-R1
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

The sample output is from the ingress router (**R1**), and shows ingress, egress, and transit LSP information. Ingress information is for the sessions that originate from this router, egress information is for sessions that terminate on this router, and transit information is for sessions that transit through this router.

There is one ingress route from **R1 (10.0.0.1)** to **R6 (10.0.0.6)**. This route is currently up, and is an active route installed in the routing table (**Rt**). The LSP **R1-to-R6** is the primary path (**P**) as opposed to the secondary path, and is indicated by an asterisk (*). The route to **R6** does not contain a named path (**ActivePath**).

There is one egress LSP from **R6** to **R1**. The **State** is up, with no routes installed in the routing table. RSVP reservation style (**Style**) consists of two parts. The first is the number of active reservations (**1**). The second is the reservation style, which is **FF** (fixed filter). The reservation style can be **FF**, **SE** (shared explicit), or **WF** (wildcard filter). There are three incoming labels (**Labelin**) and no labels going out (**Labelout**) for this LSP.

There are no transit LSPs.

For more information on checking the LSP state, see [“Checklist for Working with the Layered MPLS Troubleshooting Model” on page 1935](#).

Display Extensive Status About the LSP

Purpose

Display extensive information about LSPs, including all past state history and the reasons why an LSP might have failed.

Action

To display extensive information about LSPs, on the ingress router, enter the following Junos OS CLI operational mode command:

```
user@host> show mpls lsp extensive
```

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State:  Up , ActiveRoute:  1 , LSPname:  R1-to-R6
  ActivePath:    (primary)
  LoadBalance:  Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State:  Up
    Computed ERO  (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
```


Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.1.13.2 10.1.36.2

```

91 Aug 17 12:22:52 Selected as active path
90 Aug 17 12:22:52 Record Route: 10.1.13.2 10.1.36.2
89 Aug 17 12:22:52 Up
88 Aug 17 12:22:52 Originate Call
87 Aug 17 12:22:52 CSPF: computation result accepted
86 Aug 17 12:22:23 CSPF failed: no route toward 10.0.0.6[13920 times]
85 Aug 12 19:12:51 Clear Call
84 Aug 12 19:12:50 10.1.56.2: MPLS label allocation failure
83 Aug 12 19:12:47 Deselected as active
82 Aug 12 19:12:47 10.1.56.2: MPLS label allocation failure
81 Aug 12 19:12:47 ResvTear received
80 Aug 12 19:12:47 Down
79 Aug 12 19:12:31 10.1.56.2: MPLS label allocation failure[4 times]
78 Aug 12 19:09:58 Selected as active path
77 Aug 12 19:09:58 Record Route: 10.1.15.2 10.1.56.2
76 Aug 12 19:09:58 Up
75 Aug 12 19:09:57 Originate Call
74 Aug 12 19:09:57 CSPF: computation result accepted
73 Aug 12 19:09:29 CSPF failed: no route toward 10.0.0.6[11 times]
72 Aug 12 19:04:36 Clear Call
71 Aug 12 19:04:23 Deselected as active
70 Aug 12 19:04:23 ResvTear received
69 Aug 12 19:04:23 Down
68 Aug 12 19:04:23 CSPF failed: no route toward 10.0.0.6
67 Aug 12 19:04:23 10.1.15.2: Session preempted
66 Aug 12 16:45:35 Record Route: 10.1.15.2 10.1.56.2
65 Aug 12 16:45:35 Up
64 Aug 12 16:45:35 Clear Call
63 Aug 12 16:45:35 CSPF: computation result accepted
62 Aug 12 16:45:35 ResvTear received
61 Aug 12 16:45:35 Down
60 Aug 12 16:45:35 10.1.13.2: Session preempted
59 Aug 12 14:50:52 Selected as active path
58 Aug 12 14:50:52 Record Route: 10.1.13.2 10.1.36.2
57 Aug 12 14:50:52 Up
56 Aug 12 14:50:52 Originate Call
55 Aug 12 14:50:52 CSPF: computation result accepted
54 Aug 12 14:50:23 CSPF failed: no route toward 10.0.0.6[7 times]
53 Aug 12 14:47:22 Deselected as active
52 Aug 12 14:47:22 CSPF failed: no route toward 10.0.0.6
51 Aug 12 14:47:22 Clear Call

```



```

50 Aug 12 14:47:22 CSPF: link down/deleted
10.1.12.1(R1.00/10.0.0.1)->10.1.12.2(R2.00/10.0.0.2)
49 Aug 12 14:47:22 CSPF: link down/deleted
10.1.15.1(R1.00/10.0.0.1)->10.1.15.2(R5.00/10.0.0.5)
48 Aug 12 14:47:22 10.1.15.1: MPLS label allocation failure
47 Aug 12 14:47:22 Clear Call
46 Aug 12 14:47:22 CSPF: computation result accepted
45 Aug 12 14:47:22 10.1.12.1: MPLS label allocation failure
44 Aug 12 14:47:22 MPLS label allocation failure
43 Aug 12 14:47:22 Down
42 Jul 23 11:27:21 Selected as active path

```

Created: Sat Jul 10 18:18:44 2004

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1

From: 10.0.0.6, **LSPstate: Up** , ActiveRoute: 0

LSPname: **R6-to-R1** , LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: -

Resv style: 1 FF, Label in: 3, Label out: -

Time left: 141, Since: Tue Aug 17 12:23:14 2004

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 1 receiver 39024 protocol 0

PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 130 pkts

Adspec: received MTU 1500

PATH sentto: localclient

RESV rcvfrom: localclient

Record route: 10.1.56.2 10.1.15.2 <self>

Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Meaning

The sample output is from the ingress router (**R1**), and shows ingress, egress, and transit LSP information in detail, including all past state history and the reasons why an LSP failed. Ingress information is for sessions that originate from this router, egress information is for sessions that terminate on this router, and transit information is for sessions that transit through this router.

There is one ingress route from **R1 (10.0.0.1)** to **R6 (10.0.0.6)**. This route is currently up (**State**), with one route actively using the LSP, **R1-to-R6**. The LSP active path is the primary path. Even if the LSP does not

contain a **primary** or **secondary** keyword, the router still treats the LSP as a primary LSP, indicating that if the LSP fails, the router will attempt to signal inactive LSPs at 30-second intervals, by default.

Load balancing is **Random**, which is the default, indicating that when selecting the physical path for an LSP, the router randomly selects among equal-cost paths that have an equal hop count. Other options that you can configure are **Least-fill** and **Most-fill**. **Least-fill** places the LSP over the least utilized link of the equal-cost paths with equal hop count. **Most-fill** places the LSP over the most utilized link of the equal-cost paths sharing an equal hop count. Utilization is based on the percentage of available bandwidth.

The **Encoding type** field shows Generalized MPLS (GMPLS) signaling parameters (**Packet**), indicating IPv4. The **Switching type** is **Packet**, and the Generalized Payload Identifier (**GPID**) is IPv4.

The primary path is the active path, as indicated by an asterisk (*). The state of the LSP is **Up**.

The Explicit Route Object (**ERO**) includes the Constrained Shortest Path First (CSPF) cost (**20**) for the physical path that the LSP follows. The presence of the CSPF metric indicates that this is a CSPF LSP. The absence of the CSPF metric indicates a no-CSPF LSP.

The field **10.1.13.2 S** indicates the actual ERO. The RSVP signaling messages went to **10.1.13.2** strictly (as a next hop) and finished at **10.1.36.2** strictly. All ERO addresses are strict hops when the LSP is a CSPF LSP. Loose hops can only display in a no-CSPF LSP.

The received Record Route Object (**RRO**) has the following protection flags:

- **0x01**—Local protection available. The link downstream of this node is protected by a local repair mechanism. This flag can only be set if the Local protection flag was set in the **SESSION_ATTRIBUTE** object of the corresponding path message.
- **0x02**—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously).
- **0x04**—Bandwidth protection. The downstream router has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section.
- **0x08**—Node protection. The downstream router has a backup path providing protection against link and node failure on the corresponding path section. If the downstream router can set up only a link-protection backup path, the “Local protection available” bit is set but the “Node protection” bit is cleared.
- **0x10**—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engineered LSP. This indicates to the ingress label edge router (LER) of this LSP that it should be rerouted.

For more information on protection flags, see the *Junos Routing Protocols and Policies Command Reference*.

The field **10.1.13.2.10.1.36.2** is the actual received record route (**RRO**). Note that the addresses in the **RRO** field match those in the **ERO** field. This is the normal case for CSPF LSPs. If the RRO and ERO addresses do not match for a CSPF LSP, the LSP has to reroute or detour.

The lines numbered 91 through 42 contain the 49 most recent entries to the history log. Each line is time stamped. The most recent entries have the largest log history number and are at the top of the log, indicating that line 91 is the most recent history log entry. When you read the log, start with the oldest entry (42) to the most recent (91).

The history log was started on July 10, and displays the following sequence of activities: an LSP was selected as active, was found to be down, MPLS label allocation failed several times, was deleted several times, was preempted because of an ResvTear, was deselected as active, and was cleared. In the end, the router computed a CSPF ERO, signaled the call, the LSP came up with the listed RRO (line 90), and was listed as active.

For more information on error messages, see the *Junos MPLS Network Operations Guide Log Reference*.

The total number of ingress LSPs displayed is **1**, with **1** up and **0** down. The number in the **Up** field plus the number in the **Down** field should equal the total.

There is one egress LSP session from **R6** to **R1**. The **State** is up, with no routes installed in the routing table. RSVP reservation style (**Style**) consists of two parts. The first is the number of active reservations (**1**). The second is the reservation style, which is **FF** (fixed filter). The reservation style can be **FF**, **SE** (shared explicit), or **WF** (wildcard filter). There are three incoming labels (**Labelin**) and no labels going out (**Labelout**) for this LSP.

There are no transit LSPs.

For more information on checking the LSP state, see [“Checklist for Working with the Layered MPLS Troubleshooting Model” on page 1935](#).

Check the Status of the LSP

Purpose

Display the status of the label-switched path (LSP).

Action

To determine the LSP status, on the ingress router, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show mpls lsp
```


Sample Output

```

user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Up    1      *    R1-to-R6
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt Style  Labelin Labelout LSPname
10.0.0.1    10.0.0.6    Up    0 1 FF   3    - R6-to-R1
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output is from the ingress router (**R1**), and shows ingress, egress, and transit LSP information. Ingress information is for the sessions that originate from this router, egress information is for sessions that terminate on this router, and transit information is for sessions that transit through this router.

There is one ingress route from **R1 (10.0.0.1)** to **R6 (10.0.0.6)**. This route is currently up, and is an active route installed in the routing table (**Rt**). The LSP **R1-to-R6** is the primary path (**P**) as opposed to the secondary path, and is indicated by an asterisk (*). The route to **R6** does not contain a named path (**ActivePath**).

There is one egress LSP from **R6** to **R1**. The **State** is up, with no routes installed in the routing table. RSVP reservation style (**Style**) consists of two parts. The first is the number of active reservations (**1**). The second is the reservation style, which is **FF** (fixed filter). The reservation style can be **FF**, **SE** (shared explicit), or **WF** (wildcard filter). There are three incoming labels (**Labelin**) and no labels going out (**Labelout**) for this LSP.

There are no transit LSPs.

For more information on checking the LSP state, see [“Checklist for Working with the Layered MPLS Troubleshooting Model” on page 1935](#).

Display Extensive Status About the LSP

Purpose

Display extensive information about LSPs, including all past state history and the reasons why an LSP might have failed.

Action

To display extensive information about LSPs, on the ingress router, enter the following Junos OS CLI operational mode command:

```
user@host> show mpls lsp extensive
```

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up , ActiveRoute: 1 , LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.13.2 10.1.36.2
91 Aug 17 12:22:52 Selected as active path
90 Aug 17 12:22:52 Record Route: 10.1.13.2 10.1.36.2
89 Aug 17 12:22:52 Up
88 Aug 17 12:22:52 Originate Call
87 Aug 17 12:22:52 CSPF: computation result accepted
86 Aug 17 12:22:23 CSPF failed: no route toward 10.0.0.6[13920 times]
85 Aug 12 19:12:51 Clear Call
84 Aug 12 19:12:50 10.1.56.2: MPLS label allocation failure
83 Aug 12 19:12:47 Deselected as active
82 Aug 12 19:12:47 10.1.56.2: MPLS label allocation failure
81 Aug 12 19:12:47 ResvTear received
80 Aug 12 19:12:47 Down
79 Aug 12 19:12:31 10.1.56.2: MPLS label allocation failure[4 times]
78 Aug 12 19:09:58 Selected as active path
77 Aug 12 19:09:58 Record Route: 10.1.15.2 10.1.56.2
76 Aug 12 19:09:58 Up
75 Aug 12 19:09:57 Originate Call
74 Aug 12 19:09:57 CSPF: computation result accepted
```



```

73 Aug 12 19:09:29 CSPF failed: no route toward 10.0.0.6[11 times]
72 Aug 12 19:04:36 Clear Call
71 Aug 12 19:04:23 Deselected as active
70 Aug 12 19:04:23 ResvTear received
69 Aug 12 19:04:23 Down
68 Aug 12 19:04:23 CSPF failed: no route toward 10.0.0.6
67 Aug 12 19:04:23 10.1.15.2: Session preempted
66 Aug 12 16:45:35 Record Route: 10.1.15.2 10.1.56.2
65 Aug 12 16:45:35 Up
64 Aug 12 16:45:35 Clear Call
63 Aug 12 16:45:35 CSPF: computation result accepted
62 Aug 12 16:45:35 ResvTear received
61 Aug 12 16:45:35 Down
60 Aug 12 16:45:35 10.1.13.2: Session preempted
59 Aug 12 14:50:52 Selected as active path
58 Aug 12 14:50:52 Record Route: 10.1.13.2 10.1.36.2
57 Aug 12 14:50:52 Up
56 Aug 12 14:50:52 Originate Call
55 Aug 12 14:50:52 CSPF: computation result accepted
54 Aug 12 14:50:23 CSPF failed: no route toward 10.0.0.6[7 times]
53 Aug 12 14:47:22 Deselected as active
52 Aug 12 14:47:22 CSPF failed: no route toward 10.0.0.6
51 Aug 12 14:47:22 Clear Call
50 Aug 12 14:47:22 CSPF: link down/deleted
10.1.12.1(R1.00/10.0.0.1)->10.1.12.2(R2.00/10.0.0.2)
49 Aug 12 14:47:22 CSPF: link down/deleted
10.1.15.1(R1.00/10.0.0.1)->10.1.15.2(R5.00/10.0.0.5)
48 Aug 12 14:47:22 10.1.15.1: MPLS label allocation failure
47 Aug 12 14:47:22 Clear Call
46 Aug 12 14:47:22 CSPF: computation result accepted
45 Aug 12 14:47:22 10.1.12.1: MPLS label allocation failure
44 Aug 12 14:47:22 MPLS label allocation failure
43 Aug 12 14:47:22 Down
42 Jul 23 11:27:21 Selected as active path

```

Created: Sat Jul 10 18:18:44 2004

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1

From: 10.0.0.6, **LSPstate: Up** , ActiveRoute: 0

LSPname: **R6-to-R1** , LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: -


```

Resv style: 1 FF , Label in: 3 , Label out: -
Time left: 141, Since: Tue Aug 17 12:23:14 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39024 protocol 0
PATH rcvfrom: 10.1.15.2 (so-0/0/1.0) 130 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.56.2 10.1.15.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

The sample output is from the ingress router (**R1**), and shows ingress, egress, and transit LSP information in detail, including all past state history and the reasons why an LSP failed. Ingress information is for sessions that originate from this router, egress information is for sessions that terminate on this router, and transit information is for sessions that transit through this router.

There is one ingress route from **R1 (10.0.0.1)** to **R6 (10.0.0.6)**. This route is currently up (**State**), with one route actively using the LSP, **R1-to-R6**. The LSP active path is the primary path. Even if the LSP does not contain a **primary** or **secondary** keyword, the router still treats the LSP as a primary LSP, indicating that if the LSP fails, the router will attempt to signal inactive LSPs at 30-second intervals, by default.

Load balancing is **Random**, which is the default, indicating that when selecting the physical path for an LSP, the router randomly selects among equal-cost paths that have an equal hop count. Other options that you can configure are **Least-fill** and **Most-fill**. **Least-fill** places the LSP over the least utilized link of the equal-cost paths with equal hop count. **Most-fill** places the LSP over the most utilized link of the equal-cost paths sharing an equal hop count. Utilization is based on the percentage of available bandwidth.

The **Encoding type** field shows Generalized MPLS (GMPLS) signaling parameters (**Packet**), indicating IPv4. The **Switching type** is **Packet**, and the Generalized Payload Identifier (**GPID**) is IPv4.

The primary path is the active path, as indicated by an asterisk (*). The state of the LSP is **Up**.

The Explicit Route Object (**ERO**) includes the Constrained Shortest Path First (CSPF) cost (**20**) for the physical path that the LSP follows. The presence of the CSPF metric indicates that this is a CSPF LSP. The absence of the CSPF metric indicates a no-CSPF LSP.

The field **10.1.13.2 S** indicates the actual ERO. The RSVP signaling messages went to **10.1.13.2** strictly (as a next hop) and finished at **10.1.36.2** strictly. All ERO addresses are strict hops when the LSP is a CSPF LSP. Loose hops can only display in a no-CSPF LSP.

The received Record Route Object (**RRO**) has the following protection flags:

- **0x01**—Local protection available. The link downstream of this node is protected by a local repair mechanism. This flag can only be set if the Local protection flag was set in the `SESSION_ATTRIBUTE` object of the corresponding path message.
- **0x02**—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously).
- **0x04**—Bandwidth protection. The downstream router has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section.
- **0x08**—Node protection. The downstream router has a backup path providing protection against link and node failure on the corresponding path section. If the downstream router can set up only a link-protection backup path, the “Local protection available” bit is set but the “Node protection” bit is cleared.
- **0x10**—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engineered LSP. This indicates to the ingress label edge router (LER) of this LSP that it should be rerouted.

For more information on protection flags, see the *Junos Routing Protocols and Policies Command Reference*.

The field **10.1.13.2.10.1.36.2** is the actual received record route (**RRO**). Note that the addresses in the **RRO** field match those in the **ERO** field. This is the normal case for CSPF LSPs. If the RRO and ERO addresses do not match for a CSPF LSP, the LSP has to reroute or detour.

The lines numbered 91 through 42 contain the 49 most recent entries to the history log. Each line is time stamped. The most recent entries have the largest log history number and are at the top of the log, indicating that line 91 is the most recent history log entry. When you read the log, start with the oldest entry (**42**) to the most recent (**91**).

The history log was started on July 10, and displays the following sequence of activities: an LSP was selected as active, was found to be down, MPLS label allocation failed several times, was deleted several times, was preempted because of an ResvTear, was deselected as active, and was cleared. In the end, the router computed a CSPF ERO, signaled the call, the LSP came up with the listed RRO (line 90), and was listed as active.

For more information on error messages, see the *Junos MPLS Network Operations Guide Log Reference*.

The total number of ingress LSPs displayed is **1**, with **1** up and **0** down. The number in the **Up** field plus the number in the **Down** field should equal the total.

There is one egress LSP session from **R6** to **R1**. The **State** is up, with no routes installed in the routing table. RSVP reservation style (**Style**) consists of two parts. The first is the number of active reservations (**1**). The second is the reservation style, which is **FF** (fixed filter). The reservation style can be **FF**, **SE** (shared explicit), or **WF** (wildcard filter). There are three incoming labels (**Labelin**) and no labels going out (**Labelout**) for this LSP.

There are no transit LSPs.

For more information on checking the LSP state, see [“Checklist for Working with the Layered MPLS Troubleshooting Model” on page 1935](#).

Checking That RSVP Path Messages Are Sent and Received

Purpose

The presence or absence of various RSVP messages can help determine if there is a problem with Multiprotocol Label Switching (MPLS) in your network. For example, if path messages occur in the output without Resv messages, it might indicate that label-switched paths (LSPs) are not being created.

Action

To check that RSVP Path messages are sent and received, enter the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show rsvp statistics
```

Sample Output

```
user@R1> show rsvp statistics
```

PacketType		Total		Last 5 seconds	
		Sent	Received	Sent	Received
Path	114523	80185	1	0	
PathErr	5	10	0	0	
PathTear	12	6	0	0	
Resv FF	80515	111476	0	0	
Resv WF		0	0	0	0
Resv SE		0	0	0	0
ResvErr		0	0	0	0
ResvTear	0	5	0	0	
ResvConf		0	0	0	0
Ack		0	0	0	0
SRefresh		0	0	0	0
Hello	915851	915881	0	0	
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	0	0
Resv style conflict	0	0
Port conflict	0	0
Resv no interface	0	0
PathErr to client	15	0
ResvErr to client	0	0
Path timeout	0	0
Resv timeout	0	0
Message out-of-order	0	0
Unknown ack msg	0	0
Recv nack	0	0
Recv duplicated msg-id	0	0
No TE-link to recv Hop	0	0

Meaning

The sample output shows RSVP messages sent and received. The total number of RSVP Path messages is 11,4532 sent and 80,185 received. Within the last 5 seconds, no messages have been sent or received.

A total of 5 **PathErr** messages were sent and 10 received. When path errors occur (usually because of parameter problems in a path message), the router sends a unicast PathErr message to the sender that issued the path message. In this case, **R1** sent at least 10 path messages with an error, as indicated by the 10 PathErr messages that **R1** has received. The downstream router sent **R1** five path messages with an error, as indicated by the five PathErr messages that **R1** has sent. PathErr messages transmit in the opposite direction to path messages.

A total of 12 **PathTear** messages were sent and 6 received, none in the last 5 seconds. In contrast to PathErr messages, PathTear messages travel in the same direction as path messages. Since path messages are both sent and received, PathTear messages are also sent and received. However, if only path messages are sent, then only the PathTear messages that are sent appear in the output.

A total of 80,515 reservation (**Resv**) messages with the fixed filter (**FF**) reservation style were sent and 111,476 received, none in the last 5 seconds. An **FF** reservation style indicates that within each session, each receiver establishes its own reservation with each upstream sender, and that all selected senders are listed. No messages for the wildcard filter (**WF**) or shared explicit (**SE**) reservation styles are sent or received. For more information on RSVP reservation styles, see the *Junos MPLS Applications Configuration Guide*.

Other RSVP message types are not sent or received. For information on the ResvErr, ResvTear, and Resvconf message types, see the *Junos MPLS Applications Configuration Guide*.

Ack and summary refresh (SRefresh) messages do not appear in the output. Ack and summary refresh messages are defined in RFC 2961 and are part of the RSVP extensions. Ack messages are used to reduce the amount of RSVP control traffic in the network.

A total of 915,851 hello messages were sent and 915,881 received, with none transmitted or received in the last 5 seconds. The RSVP hello interval is 9 seconds. If more than one hello message is sent or received in the last 5 seconds, it implies that more than one interface supports RSVP.

EndtoEnd RSVP messages are legacy RSVP messages that are not used for RSVP traffic engineering. These counters increment only when RSVP forwards legacy RSVP messages issued by a virtual private network (VPN) customer for transit across the backbone to the other site(s) in the VPN. They are called end-to-end messages because they are intended for the opposite side of the network and only have meaning at the two ends of the provider network.

The **Errors** section of the output shows statistics about RSVP packets with errors. A total of 15 **PathErr to client** packets were sent to the Routing Engine. The total combines the sent and received **PathErr** packets. For more information about error statistics and packets, see the *Junos System Basics and Services Command Reference*.

Determining the Current RSVP Neighbor State

Purpose

Display a list of RSVP neighbors that were learned dynamically when exchanging RSVP packets. Once a neighbor is learned, it is never removed from the list of RSVP neighbors.

Action

To determine the current RSVP neighbor state, enter the following Junos OS CLI operational mode command:

```
user@host> show rsvp neighbor
```

Sample Output

```
user@R6> show rsvp neighbor
RSVP neighbor:  2  learned
Address      Idle Up/Dn LastChange  HelloInt  HelloTx/Rx  MsgRcvd
10.1.36.1     5  1/0  1w5d 6:30:50    9      116734/116734  23558
10.1.56.1    10  1/0  2w2d 23:44:15    9      161600/161600  23570
```


Meaning

The sample output shows that **R6** has learned about two different RSVP neighbors. Each neighbor has one line of output that includes the neighbor RSVP address, the length of time the interface was idle, the current interface up/down counter, the time of the last interface state change, the current RSVP hello interval, the total number of RSVP hello messages transmitted and received, and the total number of RSVP messages received on the interface.

The **show rsvp neighbor** command only indicates a neighbor after a session is established. Once an interface is displayed in this command output, it always appears, even if the RSVP neighbor state is down.

The RSVP neighbor **10.1.36.1** was idle for 5 seconds, came up once and has not gone down, indicating that the interface is currently in an **Up** state. As long as the up counter is one greater than the down counter, the RSVP interface is up. If the up/down counters are equal, the interface is down.

The last state change occurred 6 hours and 30 minutes ago. The current hello interval is 9 seconds. A total of 116,734 hello messages were transmitted and received on this interface, and a total of 23,558 RSVP Path/Resv messages were processed.

The RSVP neighbor **10.1.56.1** was idle for 10 seconds, came up once and has not gone down, indicating that the interface is currently in an **Up** state. The last state change occurred 23 hours and 44 minutes ago. The current Hello interval is 9 seconds. A total of 161,600 hello messages were transmitted and received on this interface, and a total of 23,570 RSVP Path/Resv messages were processed.

Take Appropriate Action

Problem

Description: Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In the example below, the routers are configured to function at different levels of the IS-IS protocol.

Solution

To correct the error in this example, enter the following commands:

Sample Output

```
[edit protocols isis]
user@R6# show
level 2 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
```



```

        disable;
    }
    interface lo0.0; {
    passive

[edit protocols isis]
user@R6# delete level 2

[edit protocols isis]
user@R6# set level 1 disable

[edit protocols isis]
user@R6# show
level 1 disable;
interface all {
    level 2 metric 10;
}
interface fxp0.0 {
    disable;
}
interface lo0.0; {
    passive

[edit protocols isis]
user@R6# commit
commit complete

[edit protocols isis]
user@R6# run show isis adjacency

```

Interface	System	L State	Hold (secs)	SNPA
so-0/0/0.0	R5	2 Up	22	
so-0/0/1.0	R4	2 Up	22	
so-0/0/2.0	R2	2 Up	22	
so-0/0/3.0	R3	2 Up	22	

Meaning

The sample output shows that the configuration error on egress router R6 has been corrected, and IS-IS adjacencies are now established.

RELATED DOCUMENTATION

Examine BGP Routes

Purpose

You can examine the BGP path selection process to determine the single, active path when BGP receives multiple routes to the same destination. In this step, we examine the reverse LSP **R6-to-R1**, making **R6** the ingress router for that LSP.

Action

To examine BGP routes and route selection, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix detail
```

Sample Output 1

```
user@R6> show route 100.100.1.1 detail
```

```
inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
100.100.1.0/24 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
              Source: 10.1.13.1
              Next hop: via so-0/0/3.0, selected
              Protocol next hop: 10.1.13.1  Indirect next hop: 8671594 304
              State: <Active Int Ext>
              Local AS: 65432 Peer AS: 65432
              Age: 4d 5:15:39      Metric2: 2
              Task: BGP_65432.10.1.13.1+3048
              Announcement bits (2): 0-KRT 4-Resolve inet.0
              AS path: I
              Localpref: 100
              Router ID: 10.0.0.1
```

Sample Output 2

```
user@R6> show route 100.100.1.1 detail
```



```

inet.0: 30 destinations, 46 routes (29 active, 0 holddown, 1 hidden)
100.100.1.0/24 (1 entry, 1 announced)
  *BGP      Preference: 170/-101
            Source: 10.0.0.1
            Next hop: via so-0/0/3.0 weight 1, selected
            Label-switched-path R6-to-R1
            Label operation: Push 100000
            Protocol next hop: 10.0.0.1   Indirect next hop: 8671330 301
            State: <Active Int Ext>
            Local AS: 65432 Peer AS: 65432
            Age: 24:35      Metric2: 2
            Task: BGP_65432.10.0.0.1+179
            Announcement bits (2): 0-KRT 4-Resolve inet.0
            AS path: I
            Localpref: 100
            Router ID: 10.0.0.1

```

Meaning

Sample Output 1 shows that the BGP next hop (**10.1.13.1**) does not equal the LSP destination address (**10.0.0.1**) in the **to** statement at the `[edit protocols mpls label-switched-path label-switched-path-name]` hierarchy level when the BGP configuration of **R6** and **R1** is incorrect.

Sample Output 2, taken after the configurations on R1 and R6 are corrected, shows that the BGP next hop (**10.0.0.1**) and the LSP destination address (**10.0.0.1**) are the same, indicating that BGP can use the LSP to forward BGP traffic.

CLI Operational Mode Top-Level Commands

In operational mode, you enter commands to monitor and diagnose the software, network connectivity, and the router. When you log in to the router and the CLI starts, you are at the top level of the CLI operational mode. At this level, there are several broad groups of CLI commands. [Table 65 on page 2304](#) lists the top-level CLI operational mode commands and describes the options available for each command. The commands are listed in alphabetical order.

Table 65: CLI Operational Mode Top-Level Commands

Command	Description
clear	<p>Clear statistics and protocol database information.</p> <p>Syntax: clear (arp bgp firewall helper igmp ike ilmi interfaces ipsec ipv6 isis ldp log mpls msdp multicast ospf pim rip ripng route rsvp snmp system vrrp)</p>
configure	<p>Enter CLI configuration mode.</p> <p>Alternative commands: configure <exclusive> <private></p>
file	<p>Perform file manipulation operations, such as copy, delete, list, rename, and show.</p> <p>Syntax: file (compare copy delete list rename show)</p>
help	<p>Provide help information.</p> <p>Syntax: help (reference syslog topic)</p>
monitor	<p>Monitor a log file or interface traffic in real time.</p> <p>Syntax: monitor (interface list start stop traffic)</p>
mtrace	<p>Display trace information about a multicast path from a source to a receiver.</p> <p>Syntax: mtrace (from-source monitor to-gateway)</p>
ping	<p>Verify IP connectivity to another IP host or Asynchronous Transfer Mode (ATM) connectivity (ping ATM) using Operation Administration and Maintenance (OAM) cells to an ATM endstation.</p> <p>Syntax: ping host <interface source-interface > <bypass-routing> <count requests > <do-not-fragment> <interval seconds > <pattern string > <record-route> <routing-instance routing-instance-name > <size bytes > <strict> <tos type-of-service > <tll value > <via route > <rapid detail></p> <p>Syntax: ping atm interface interface <count count > <end-to-end segment> <interval interval> <sequence-number sequence-number > <vci vci > <brief></p> <p>Syntax: ping vpn-interface vpn-interface host <local echo-address></p>
pipe	<p>Filter the output of an operational mode or configuration mode command.</p> <p>Syntax: (compare count display <detail inheritance xml> except <i>pattern</i> find <i>pattern</i> last <i>lines</i> match <i>pattern</i> no-more resolve <file-names> save <i>filename</i> trim <i>columns</i>)</p>
quit	<p>Log out from the CLI process.</p> <p>Syntax: quit</p>

Table 65: CLI Operational Mode Top-Level Commands (*continued*)

Command	Description
request	<p>Make system-level requests, such as halt or reboot the router, load software packages, and back up the router's file systems.</p> <p>Syntax: request system (halt reboot snapshot software)</p>
restart	<p>Restart the router hardware or software processes.</p> <p>Syntax: restart (fpc class-of-service gracefully immediately interface-control mib-process network-access-service remote-operations routing sampling sfm snmp soft)</p>
set	<p>Set CLI properties, the router's date and time, and the craft interface display text.</p> <p>Syntax: set (chassis cli date)</p>
show	<p>Show information about all aspects of the software, including interfaces and routing protocols.</p> <p>Syntax: show (accounting aps arp as-path bgp chassis cli configuration connections dvmrp firewall helper host igmp ike ilmi interfaces ipsec ipv6 isis l2circuit l2vpn ldp link-management log mpls msdp multicast ntp ospf pfe pim policer policy rip ripng route rsvp sap snmp system task ted version vrrp)</p>
ssh	<p>Open a secure shell to another host.</p> <p>Syntax: ssh host <bypass-routing> <routing-instance routing-instance-name> <source address> <vpn-interface vpn-interface> <v1 v2></p>
start	<p>Start a software process.</p> <p>Syntax: start shell</p>
telnet	<p>Start a telnet session to another host.</p> <p>Syntax: telnet host <8bit> <bypass-routing> <inet inet6> <noresolve> <port port> <interface interface-name> <routing-instance routing-instance-name> <source address> <vpn-interface vpn-interface></p>
test	<p>Run various diagnostic debugging commands.</p> <p>Syntax: test (configuration interface msdp policy)</p>
traceroute	<p>Trace the route to a remote host.</p> <p>Syntax: traceroute host <as-number-lookup> <bypass-routing> <gateway address> <inet inet6> <noresolve> <routing-instance routing-instance-name> <source address> <tos value> <ttd value> <vpn-interface vpn-interface> <wait seconds></p>

CLI Keyboard Shortcuts

In the CLI, you can use keyboard sequences to move around and edit a command line. You can also use keyboard sequences to scroll through a list of recently executed commands.

Table 15 on page 23 lists some of the CLI keyboard sequences.

Table 66: CLI Keyboard Shortcuts

Keyboard sequence	Action
Ctrl+b	Move the cursor back one character.
Esc+b or Alt+b	Move the cursor back one word.
Ctrl+f	Move the cursor forward one character.
Esc+f or Alt+f	Move the cursor forward one word.
Ctrl+a	Move the cursor to the beginning of the command line.
Ctrl+e	Move the cursor to the end of the command line.
Ctrl+h, Delete, or Backspace	Delete the character before the cursor.
Ctrl+d	Delete the character at the cursor.
Ctrl+k	Delete the all characters from the cursor to the end of the command line.
Ctrl+u or Ctrl+x	Delete the all characters from the command line.
Ctrl+w, Esc+Backspace, or Alt+Backspace	Delete the word before the cursor.
Esc+d or Alt+d	Delete the word after the cursor.
Ctrl+y	Insert the most recently deleted text at the cursor.
Ctrl+l	Redraw the current line.
Ctrl+p	Scroll backward through the list of recently executed commands.
Ctrl+n	Scroll forward through the list of recently executed commands.

Table 66: CLI Keyboard Shortcuts (*continued*)

Keyboard sequence	Action
Ctrl+r	Search the CLI history incrementally in reverse order for lines matching the search string.
Esc+/ or Alt+/ 	Search the CLI history for words for which the current word is a prefix.
Esc+. or Alt+ 	Scroll backward through the list of recently entered words in a command line.
Esc+ <i>number sequence</i> or Alt+ <i>number sequence</i>	Specify the number of times to execute a keyboard sequence.

Manage Output at the ---(more)--- Prompt

If the output is longer than the screen length, it appears one screen at a time with the UNIX ---(more)--- prompt at the end of the screen. The ---(more)--- prompt indicates that more output is available. The following table lists the keyboard sequences you can use at the --(more)--- prompt.

Table 67: Keyboard Shortcuts at the ---(more)--- Prompt

Keyboard Shortcut	Action
Enter, Return, k, Ctrl+m, Ctrl+n, or down arrow	Scroll down one line.
Tab, d, Ctrl+d, or Ctrl+x	Scroll down one-half screen.
Space or Ctrl+f	Scroll down one whole screen.
Ctrl+e or g	Scroll down to the bottom of the output.
n (or no-more)	Display the output all at once instead of one screen at a time.
j, Ctrl-h, Ctrl-p, or up arrow	Scroll up one line.
u or Ctrl-u	Scroll up one-half screen.
b or Ctrl-b	Scroll up one whole screen.
Ctrl-a or g	Scroll up to the bottom of the output.
/string	Search forward for a string.

Table 67: Keyboard Shortcuts at the ---(more)--- Prompt (*continued*)

Keyboard Shortcut	Action
? <i>string</i>	Search backward for a string.
n	Repeat previous search for a string.
m or M (or <i>match string</i>)	Find a text string. You are prompted for the string to match
e or E (or <i>except string</i>)	Find, ignoring a text string. You are prompted for the string to ignore.
Ctrl-C, q, Q, or Ctrl-k	Interrupt the display of output.
H (Same as specifying <i>hold</i>)	Hold the CLI at the —more- prompt after displaying all output.
c or C	Clear any match conditions and display the complete output.
Ctrl-l	Redraw the output on the screen.
s or S (or <i>save filename</i>)	Save the command output to a file. You are prompted for a filename.

Working with Problems on Your Network

Problem

Description: This checklist provides links to troubleshooting basics, an example network, and includes a summary of the commands you might use to diagnose problems with the router and network.

Solution

Table 68: Checklist for Working with Problems on Your Network

Tasks	Command or Action
“Isolating a Broken Network Connection” on page 2309	
1. <i>Identifying the Symptoms of a Broken Network Connection</i>	<p>ping (<i>ip-address</i> <i>hostname</i>)</p> <p>show route (<i>ip-address</i> <i>hostname</i>)</p> <p>traceroute (<i>ip-address</i> <i>hostname</i>)</p>
2. <i>Isolating the Causes of a Network Problem</i>	<p>show < configuration interfaces protocols route</p> <p>></p>


Table 68: Checklist for Working with Problems on Your Network (continued)

Tasks	Command or Action
3. Taking Appropriate Action for Resolving the Network Problem	<code>[edit]</code> <code>delete routing options static route destination-prefix</code> <code>commit and-quit</code> <code>show route destination-prefix</code>
4. Evaluating the Solution to Check Whether the Network Problem Is Resolved	<code>show route (ip-address hostname)</code> <code>ping (ip-address hostname) count 3</code> <code>traceroute (ip-address hostname)</code>

Isolating a Broken Network Connection

By applying the standard four-step process illustrated in [Figure 171 on page 2309](#), you can isolate a failed node in the network. Note that the functionality described in this section is not supported in versions 15.1X49, 15.1X49-D30, or 15.1X49-D40.


Figure 171: Process for Diagnosing Problems in Your Network



Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic in your approach to the maintenance and monitoring of your network. To prepare for problems on your network, understand how the network functions under normal conditions, have records of baseline network activity, and carefully observe the behavior of your network during a problem situation.

[Figure 172 on page 2310](#) shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

Figure 172: Network with a Problem



The network in [Figure 172 on page 2310](#) consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (**R1**) in AS 65001 announces aggregated prefixes **100.100/24** to the AS 65002 network. The problem in this network is that **R6** does not have access to **R5** because of a loop between **R2** and **R6**.

To isolate a failed connection in your network, follow the steps in these topics:

- *Isolating the Causes of a Network Problem*
- *Taking Appropriate Action for Resolving the Network Problem*
- *Taking Appropriate Action for Resolving the Network Problem*
- *Evaluating the Solution to Check Whether the Network Problem Is Resolved*

Display Junos OS Information

Purpose

Display Junos OS information and status to determine if the version of Junos OS that you are running supports particular features or hardware. You can also determine whether particular software bugs will affect your version of Junos OS.

Action

To display Junos OS information, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show version
```

Sample Output

```
user@host> show version
```

```
Hostname: my-router.net
Model: m160
JUNOS Base OS boot [5.5R2.3]
JUNOS Base OS Software Suite [5.5R2.3]
JUNOS Kernel Software Suite [5.5R2.3]
JUNOS Packet Forwarding Engine Support [5.5R2.3]
JUNOS Routing Software Suite [5.5R2.3]
JUNOS Online Documentation [5.5R2.3]
JUNOS Crypto Software Suite [5.5R2.3]
KERNEL 5.5R2.3 #0 built by builder on 2002-11-21 22:56:20 UTC
MGD release 5.5R2.3 built by builder on 2002-11-21 22:36:05 UTC
CLI release 5.5R2.3 built by builder on 2002-11-21 22:33:44 UTC
CHASSISD release 5.5R2.3 built by builder on 2002-11-21 22:32:10 UTC
DCD release 5.5R2.3 built by builder on 2002-11-21 22:30:06 UTC
RPD release 5.5R2.3 built by builder on 2002-11-21 22:37:08 UTC
SNMPD release 5.5R2.3 built by builder on 2002-11-21 22:43:14 UTC
MIB2D release 5.5R2.3 built by builder on 2002-11-21 22:36:10 UTC
APSD release 5.5R2.3 built by builder on 2002-11-21 22:32:07 UTC
VRRPD release 5.5R2.3 built by builder on 2002-11-21 22:43:26 UTC
ALARMD release 5.5R2.3 built by builder on 2002-11-21 22:32:01 UTC
PFED release 5.5R2.3 built by builder on 2002-11-21 22:36:53 UTC
CRAFTD release 5.5R2.3 built by builder on 2002-11-21 22:33:59 UTC
SAMPLED release 5.5R2.3 built by builder on 2002-11-21 22:43:01 UTC
ILMID release 5.5R2.3 built by builder on 2002-11-21 22:35:17 UTC
RMOPD release 5.5R2.3 built by builder on 2002-11-21 22:37:01 UTC
COSD release 5.5R2.3 built by builder on 2002-11-21 22:33:50 UTC
KMD release 5.5R2.3 built by builder on 2002-11-21 22:35:29 UTC
FSAD release 5.5R2.3 built by builder on 2002-11-21 22:34:14 UTC
SERVICED release 5.5R2.3 built by builder on 2002-11-21 22:43:07 UTC
IRSD release 5.5R2.3 built by builder on 2002-11-21 22:35:21 UTC
NASD release 5.5R2.3 built by builder on 2002-11-21 22:36:47 UTC
FUD release 5.5R2.3 built by builder on 2002-11-21 22:34:17 UTC
PPMD release 5.5R2.3 built by builder on 2002-11-21 22:36:58 UTC
LMPD release 5.5R2.3 built by builder on 2002-11-21 22:36:01 UTC
RTSPD release 5.5R2.3 built by builder on 2002-11-21 22:42:58 UTC
```



```
SMARTD release 5.5R2.3 built by builder on 2002-11-21 22:47:50 UTC
jkernel-dd release 5.5R2.3 built by builder on 2002-11-21 22:27:20 UTC
jroute-dd release 5.5R2.3 built by builder on 2002-11-21 22:27:34 UTC
jcrypto-dd release 5.5R2.3 built by builder on 2002-11-21 22:27:46 UTC
```

Meaning

The sample output shows the hostname, the version information for the Junos OS packages installed on the machine, and the version information for each software process.

Display Version Information for Junos OS Packages

Purpose

Display version information for Junos OS packages to determine if they support particular features or hardware. You can also determine whether particular software bugs will affect your version of Junos OS.

Action

To display brief information and status for the kernel and Packet Forwarding Engine, use the following CLI operational mode command:

```
user@host> show version brief
```

The following sample output is for worldwide nonencrypted Junos OS:

Sample Output

```
user@host> show version brief
Hostname: my-router.net
Model: m10
JUNOS Base OS boot [5.5R2.3]
JUNOS Base OS Software Suite [5.5R2.3]
JUNOS Kernel Software Suite [5.5R2.3]
JUNOS Packet Forwarding Engine Support [5.5R2.3]
JUNOS Routing Software Suite [5.5R2.3]
JUNOS Online Documentation [5.5R2.3]
```

The following sample output is for Canada and USA encrypted Junos OS:

```
user@host> show version brief
Hostname: my-router.net
```



```

Model: m10
JUNOS Base OS boot [5.5R2.3]
JUNOS Base OS Software Suite [5.5R2.3]
JUNOS Kernel Software Suite [5.5R2.3]
JUNOS Packet Forwarding Engine Support [5.5R2.3]
JUNOS Routing Software Suite [5.5R2.3]
JUNOS Online Documentation [5.5R2.3]
JUNOS Crypto Software Suite [5.5R2.3]

```

Meaning

The sample output shows version information for the Junos OS packages installed on the router. If the **Junos Crypto Software Suite** is listed, the router has Canada and USA encrypted Junos OS. If the **Junos Crypto Software Suite** is not listed, the router is running worldwide nonencrypted Junos OS.

Display the Current Active Router Configuration

Purpose

Examine the current active router configuration.

Action

To display the current, active router configuration, use the following command-line interface (CLI) operational mode command:

```
user@host> show configuration
```

Sample Output

```

user@host> show configuration
version "10.4R2";
groups {
  global {
    system {
      host-name potter;
      domain-name harry.potter.net;
      domain-search [ harry.potter.net potter.net hrypтр.net ];
      backup-router 10.110.12.254;
      time-zone America/Los_Angeles;
    }
  }
}

```



```

authentication-order [ tacplus
password radius ];
    root-authentication {
        encrypted-password "$1$Off5.$I7.kUgMmx/4WKwUAG"; # SECRET-DATA
    }
    name-server {
        172.17.28.101;
        172.17.28.100;
    }
    radius-server {
        10.168.5.73 {
            secret "$9$Nd-YoDjq.PT4oZjik5T369pBIhSlL7dC"; # SECRET-DATA
            timeout 5;
            retry 3;
        }
    }
    tacplus-server {
        10.168.5.73 {
            secret "$9$.539IRSM87O1lMX-2gqmfTz6"; # SECRET-DATA
            timeout 15;
            single-connection;
        }
    }
    login {
        class superuser-local {
            permissions all;
        }
        class wheel {
            permissions [ admin clear field floppy interfacemaintenance
network reset routing shell snmp system trace view ];
        }
        class readonly {
            permissions [ interface network routing system trace view ];
        }
        user rpe {
            uid 1230;
            class superuser;
            shell csh;
            authentication {
                encrypted-password FN5oyk/qZO7F2; # SECRET-DATA
            }
        }
        [...]
    }
}

```



```

    }
    static-host-mapping {
        crater sysid 0102.5524.5045;
        badlands sysid 0102.5524.5046;
        [...Output truncated...]
    }
    services {
        finger;
        ftp;
        rlogin;
        rsh;
        ssh;
telnet;
    }
    syslog {
        user * {
            any emergency;
        }
        host log {
            any notice;
            pfe info;
            interactive-commands any;
        }
        file messages {
            any notice;
            authorization info;
            pfe info;
            archive world-readable;
        }
        file security {
            interactive-commands any;
            archive world-readable;
        }
        file white_bx {
            daemon notice;
            archive size 40m world-readable;
        }
    }
    processes {
        routing enable;
        snmp enable;
        tnp-process enable;
        ntp enable;
        inet-process enable;
    }

```



```

        mib-process enable;
        management enable;
        watchdog enable;
    }
    ntp {
        boot-server ntp.juniper.net;
        server 172.17.27.46;
    }
}
chassis {
    dump-on-panic;
}
snmp {
    location "Systest lab";
    contact "Brian Matheson";
    interface fxp0.0;
    community public {
        authorization read-only;
    }
    community private {
        authorization read-write;
    }
}
routing-options {
    static {
        /* corporate and alpha net */
        route 172.16.0.0/12 {
            next-hop 10.168.14.254;
            retain;
            no-readvertise; [...Output truncated...]
        }
    }
}
}
rel;
}

apply-groups [ global re0 re1 ];
chassis {
    fpc 0 {
        pic 0 {
            mlfr-uni-nni-bundles 4;
        }
    }
}

```



```

}
interfaces {
    ls-0/0/0:0 {
        encapsulation multilink-frame-relay-uni-nni;
        unit 0 {
            dlci 100;
            family inet {
                address 10.53.99.2/32 {
                    destination 10.53.99.1;
                }
            }
        }
    }
    ct3-0/1/0 {
        partition 1 interface-type t1;
        partition 2 interface-type t1;
        partition 3 interface-type t1;
        partition 4 interface-type t1;
    }
    t1-0/1/0:1 {
        encapsulation multilink-frame-relay-uni-nni;
        unit 0 {
            family mlfr-uni-nni {
                bundle ls-0/0/0:0;
            }
        }
    }
}
routing-options {
    static {
        route 10.1.1.0/24 next-hop 10.53.99.1;
    }
    autonomous-system 69;
    forwarding-table {
        export pplb;
    }
}
protocols {
    bgp {
        disable;
        group int {
            type internal;
            neighbor 10.255.14.30;
            [...Output truncated...] }
        }
    }
}

```



```

isis {
    disable;
    interface all {
        level 1 disable;
    }
    interface fxp0.0 {
        disable;
    }
}
inactive: ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
}
policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}
[...Output truncated...]

```

Meaning

The sample output shows the current, active configuration for the router. When displaying the configuration, the CLI indents each subordinate hierarchy level, inserts braces to indicate the beginning and end of each hierarchy level, and places semicolons at the end of statements that are at the lowest level of the hierarchy.

The configuration statements appear in a fixed order. Interfaces appear alphabetically by type, and then in numerical order by slot number, Physical Interface Card (PIC) number, and port number.

Copy Junos OS to the Router

Action

Copy the software packages from the server to the router. We recommend that you copy them to the **/var/tmp** directory, which is on the rotating medium (hard disk) and is a large file system. Use the following CLI command:

```
user@host> file copy ftp://username: prompt@ftp.hostname.
net/jinstall-package-name/var/tmp/jinstall-package-name
```

Add New Software

Action

To add new software packages, use the following Junos OS CLI operational mode command:

```
user@host> request system software add /var/tmp/jinstall-package--name
```

package-name is the full URL to the file and **release-number** is the major software release number; for example, 4.2R1. Before the new software is added, the existing software is automatically deleted.

NOTE: Even though you are adding the new software, the changes do not take effect until the router has completed rebooting.

Sample Output

```
user@host> request system software add /var/tmp/jinstall-5.2R2.3-domestic.tgz
Installing package '/var/tmp/jinstall-5.2R2.3-domestic.tgz'
Auto-deleting old jroute...
Auto-deleting old jdocs...
Auto-deleting old jpfe...
Auto-deleting old jkernel...
Adding JUNOS base software 5.2R2.3
Adding jkernel...
Adding jpfe...
Adding jdocs...
Adding jroute...
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Saving package file in /var/sw/pkg/jinstall-5.2R2.3-domestic.tgz
```


Compare Information Logged Before and After the Upgrade

Purpose

Compare the operation of the system before and after the upgrade to ensure that everything is working as expected.

Action

To obtain system information, use the following Junos OS CLI operational mode commands:

```
user@host> show version
user@host> show chassis hardware
user@host> show configuration
user@host> show interface terse
user@host> show bgp summary
user@host> show isis adjacency brief
user@host> show ospf neighbor brief
user@host> show system storage
```

Compare the information from these commands with the information you logged before the upgrade.

Displaying LSP Status Events

Purpose

Display extensive information about LSPs, including the 50 most recent history events and the reasons why an LSP might have failed.

Action

To examine status messages, enter the following Junos OS command-line interface (CLI) operational mode command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1# run show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
```



```

From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
ActivePath: (none)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary                               State: Dn
    Will be enqueued for recomputation in 3 second(s).
68 Jan  5 10:02:56 CSPF failed: no route toward 10.0.0.6[9 times]
67 Jan  5 09:58:33 Deselected as active
66 Jan  5 09:58:33 CSPF failed: no route toward 10.0.0.6
65 Jan  5 09:58:33 Clear Call
64 Jan  5 09:58:33 Session preempted
63 Jan  5 09:58:33 Down
62 Jan  5 09:58:32 CSPF failed: no route toward 10.0.0.6[2 times]
61 Jan  5 09:57:55 10.1.36.2: Explicit Route: wrong delivery
60 Jan  5 09:57:34 CSPF failed: no route toward 10.0.0.6[2 times]
59 Jan  5 09:57:28 CSPF: link down/deleted
10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)
58 Jan  5 09:54:37 Selected as active path
57 Jan  5 09:54:37 Record Route: 10.1.13.2 10.1.36.2
56 Jan  5 09:54:37 Up
55 Jan  5 09:54:37 Originate Call
54 Jan  5 09:54:37 CSPF: computation result accepted
53 Jan  4 18:11:28 CSPF failed: no route toward 10.0.0.6[2 times]
52 Jan  4 18:10:44 Deselected as active
51 Jan  4 18:10:44 CSPF failed: no route toward 10.0.0.6
50 Jan  4 18:10:44 CSPF: link down/deleted
10.1.13.1(R1.00/10.0.0.1)->10.1.13.2(R3.00/10.0.0.3)
49 Jan  4 18:10:44 RSVP Disabled
48 Jan  4 18:10:44 RSVP error , subcode 4: protocol shutdown
47 Jan  4 18:10:44 Down
46 Jan  4 18:06:15 Up
45 Jan  4 18:06:15 Down
44 Jan  4 18:06:10 Selected as active path
43 Jan  4 18:06:09 Record Route: 10.1.13.2 10.1.36.2
42 Jan  4 18:06:09 Up
41 Jan  4 18:06:09 Originate Call
40 Jan  4 18:06:09 CSPF: computation result accepted
39 Jan  4 18:05:40 CSPF failed: no route toward 10.0.0.6[2 times]
38 Jan  4 18:04:57 Deselected as active
37 Jan  4 18:04:57 CSPF failed: no route toward 10.0.0.6
36 Jan  4 18:04:57 CSPF: link down/deleted
10.1.13.1(R1.00/10.0.0.1)->10.1.13.2(R3.00/10.0.0.3)
35 Jan  4 18:04:57 CSPF failed: no route toward 10.0.0.6

```



```

34 Jan  4 18:04:57 Clear Call
33 Jan  4 18:04:57 Explicit Route: bad strict route
32 Jan  4 18:04:57 No Route toward dest
31 Jan  4 18:04:57 Down
30 Dec 28 13:47:29 Selected as active path
29 Dec 28 13:47:29 Record Route: 10.1.13.2 10.1.36.2
28 Dec 28 13:47:29 Up
27 Dec 28 13:47:29 Originate Call
26 Dec 28 13:47:29 CSPF: computation result accepted
25 Dec 28 13:46:59 CSPF failed: no route toward 10.0.0.6
24 Dec 28 13:46:39 Deselected as active
23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6
22 Dec 28 13:46:39 Clear Call
21 Dec 28 13:46:39  ResvTear received
20 Dec 28 13:46:39 Down
19 Dec 28 13:46:39 10.1.13.2:  Session preempted
Created: Mon Dec 13 11:47:18 2004
Total 1 displayed, Up 0, Down 1
[...Output truncated...]

```

Sample Output 2

```

user@R1> show mpls lsp extensive
[...Output truncated...]
*Primary use-TOKYO      State: Up, No-decrement-ttl

  Received RRO:
    10.222.28.2(flag=0x9) 10.222.4.2(flag=0x1) 10.222.44.2
    7 Sep 20 18:13:45 Record Route: 10.222.28.2(flag=0x9)
    10.222.4.2(flag=0x1) 10.222.44.2
    6 Sep 20 18:13:45 Record Route: 10.222.28.2(flag=0x9)
    10.222.4.2 10.222.44.2
    5 Sep 20 18:13:45 Fast-reroute Detour Up
    4 Sep 20 18:13:42 Selected as active path
    3 Sep 20 18:13:42 Record Route: 10.222.28.2 10.222.4.2
    10.222.44.2
    2 Sep 20 18:13:42 Up
    1 Sep 20 18:13:42 Originate Call

```


Sample Output 3

```

user@R1> show mpls lsp extensive
[...Output truncated...]
*Primary    long                State: Up, COS: 6
  Bandwidth per class: <ct0 20Mbps> <ct1 2Mbps> <ct2 3Mbps>
  OptimizeTimer: 250
  Reoptimization in 237 second(s).

  Computed ERO
(S [L] denotes strict [loose] hops): (CSPF metric: 50)
    10.35.38.2 S 192.168.135.29
S 10.35.39.1 S 10.35.40.2 S 10.35.41.1 S

  Received RRO
(ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
    10.35.38.2 (flag=0x09) 192.168.135.29 (flag=0x10) 10.35.39.1
(flag=0x01) 10.35.40.2 (flag=0x01) 10.35.41.1 (flag=0x01)
[...Output truncated...]

```

Meaning

Sample Output 1 from ingress router R1 shows extensive ingress LSP information, including LSP events that led to an LSP failure and the 50 most recent state events.

LSP events in bold are described in this topic. Descriptions include sample output of the LSP event, an explanation of what the event means, the possible cause of the event, and any possible actions that you can take.

For completeness, events not included in this example output are also described in this topic to show LSP events that did not occur in the example network configuration, but might occur in your network. The LSP events are organized alphabetically.

Sample Output 2 shows the state of the route received in the Received Record Route (**Received RRO**) created by fast reroute configurations in the network. The **Received RRO** indicates a series of hops. Each hop has an address followed by a flag. For more information on flags, see the *Junos MPLS Network Operations Guide*. In most cases, the **Received RRO** is the same as the computed Explicit Route Object (ERO).

Sample Output 3 shows a **Computed ERO** and a **Received RRO**. In this instance they are the same. However, if **Received RRO** is different from the **Computed ERO**, there is a topology change in the network, and the route is taking a detour.

Call Was Cleared by RSVP Event

LSP Event

Call was cleared by RSVP

Sample Output

```
user@R1> show mpls lsp extensive
[...Output truncated...]
10.0.0.6
From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
ActivePath: (none)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary State: Dn
Will be enqueued for recomputation in 10 second(s).
11 Jan 26 14:58:32 CSPF failed: no route toward 10.0.0.6
10 Jan 26 14:58:25 Deselected as active
9 Jan 26 14:58:25 CSPF failed: no route toward 10.0.0.6
8 Jan 26 14:58:25 Call was cleared by RSVP
7 Jan 26 14:58:25 Session preempted
6 Jan 26 14:58:25 Down
[...Output truncated...]
```

Meaning

This LSP event indicates that the Resource Reservation Protocol (RSVP) session corresponding to the LSP path was preempted and the corresponding RSVP state deleted.

Cause

This LSP event occurs when you issue the **clear rsvp session** command or trigger preemption of an RSVP session at the ingress router. Depending on the timer value, Constrained Shortest Path First (CSPF) recomputes the path and the LSP comes up again.

Action

Not applicable.

Change in Active Path Event

LSP Event

Change in active path

Sample Output

```

user@R1> show mpls lsp extensive
[...Output truncated...]
13 Sep 19 00:02:20 Deselected as active
12 Sep 19 00:02:20 ResvTear received
11 Sep 19 00:02:20 Down
10 Sep 19 00:02:20   Change in active path
9 Sep 19 00:02:20
8 Sep 19 00:02:20 10.222.28.2: Explicit Route: bad strict routeChange in active
path
7 Sep 19 00:02:20 CSPF failed: no route toward 192.168.32.1
6 Sep 19 00:02:20 10.222.28.2: No Route toward dest
5 Sep 19 00:00:54 Selected as active path
4 Sep 19 00:00:54 Record Route: 10.222.28.2 10.222.4.2 10.222.44.2
3 Sep 19 00:00:54 Up
2 Sep 19 00:00:54 Originate Call
1 Sep 19 00:00:54 CSPF: computation result accepted
[...Output truncated...]

```

Meaning

This LSP event indicates that even though the active physical path has changed, the LSP stays up. Because this network configuration has an alternate (fast-reroute) path available, the event is a **Change in active path** rather than a **Session preempted** event.

Cause

The active path might have failed.

Action

Not applicable.

Clear Call Event

LSP Event

Clear call

Sample Output

```

user@R1> show mpls lsp extensive
[...Output truncated...]

```



```

65 Jan  5 09:58:33  Clear Call
64 Jan  5 09:58:33 Session preempted
63 Jan  5 09:58:33 Down
[...Output truncated...]

```

Meaning

This LSP event indicates that the LSP was disconnected and restarted.

Cause

The **clear mpls lsp** command was issued on the ingress router to disconnect existing RSVP sessions, release the routes and states associated with the LSP, and then start a new LSP. Issuing this command might impact traffic travelling along the LSP, because a time lag might occur between tearing down the old path and setting up a new path.

Action

Not applicable.

Deselected as Active Event

LSP Event

Deselected as active

Sample Output

```

user@R1> show mpls lsp extensive
[...Output truncated...]
Will be enqueued for recomputation in 18 second(s).
53 Jan  4 18:11:28 CSPF failed: no route toward 10.0.0.6[2 times]
52 Jan  4 18:10:44  Deselected as active
51 Jan  4 18:10:44 CSPF failed: no route toward 10.0.0.6
50 Jan  4 18:10:44 CSPF: link down/deleted
[...Output truncated...]

```

Meaning

This LSP event indicates that the LSP is no longer the active path.

Cause

Typically, other events, similar to those in lines 50 and 51, indicate the reason that the LSP is no longer the active path.

Action

Refer to events on either side of this event to determine the appropriate action.

Link Protection Down Event

LSP Event

Link protection down

Sample Output 1

```
user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
    link-protection;
}
interface fxp0.0 {
    disable;
}
interface all;
```

Sample Output 2

```
user@R1> show mpls lsp extensive
[...Output truncated...]
10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
10.1.13.2 S 10.1.36.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.13.2(flag=1 Label=101936) 10.1.36.2(Label=3)
70 Feb 10 11:01:56 Link-protection Up
69 Feb 10 11:01:56 Selected as active path
68 Feb 10 11:01:56 Link-protection Down
  67 Feb 10 11:01:56 Link-protection Up
```



```

66 Feb 10 11:01:56 Record Route: 10.1.13.2(flag=1 Label=101936)
10.1.36.2(Label=3)
65 Feb 10 11:01:56 Up
64 Feb 10 11:01:56 Originate Call
63 Feb 10 11:01:56 CSPF: computation result accepted
62 Feb 10 11:01:56 Clear Call
61 Feb 10 11:01:56 Deselected as active
60 Feb 10 11:01:56 Link-protection Down
59 Feb 10 10:57:58 Record Route: 10.1.13.2(flag=1 Label=101920)
10.1.36.2(Label=3)
58 Feb 10 10:57:56 Link-protection Up
57 Feb 10 10:56:58 Selected as active path
56 Feb 10 10:56:58 Record Route: 10.1.13.2(Label=101920) 10.1.36.2(Label=3)
55 Feb 10 10:56:58 Up
54 Feb 10 10:56:58 Originate Call
53 Feb 10 10:56:58 CSPF: computation result accepted
52 Feb 10 10:56:58 Clear Call
51 Feb 10 10:56:58 Deselected as active
50 Feb 10 10:56:58 Link-protection Down
49 Feb 10 10:56:35 10.1.56.2: MPLS label allocation failure[2 times]
48 Feb 10 10:48:32 Link-protection Up
47 Feb 10 10:48:32 Selected as active path
[...Output truncated...]

```

Meaning

Sample Output 1 shows the MPLS link-protection configuration on R1 for the LSP **R1-to-R6**.

Sample Output 2 shows that link protection came up and down several times. Link protection comes up when the LSP signals. Line 60 shows the result when RSVP is disabled on all alternate paths out of R6. Lines 68 to 70 are the result when the **clear mpls lsp** command is issued.

Cause

This LSP event is caused by a failure or configuration change that deletes or resignals the bypass LSP. For example, you clear the LSP using the **clear mpls lsp** command, or you disable RSVP on all alternate paths for the LSP. The bypass LSP does not use the primary path, instead it looks for an alternate path.

Action

Include the **family mpls** statement for all alternate paths for the LSP at the [edit interfaces type-fpc/pic/port.unit] hierarchy level.

Originate Call Event

LSP Event

Originate call

Sample Output

```
user@R1> show mpls lsp extensive
[...Output truncated...]
 43 Jan  4 18:06:09 Record Route:  10.1.13.2 10.1.36.2
 42 Jan  4 18:06:09 Up
 41 Jan  4 18:06:09  Originate Call
 40 Jan  4 18:06:09 CSPF: computation result accepted
 39 Jan  4 18:05:40 CSPF failed: no route toward 10.0.0.6[2 times]
[...Output truncated...]
```

Meaning

This LSP event indicates that the router is issuing an RSVP Path message.

Cause

A Path message is transmitted by the ingress router toward the egress router to establish an LSP.

Action

To analyze the contents of the Path message, enable RSVP tracing. To configure RSVP tracing, include the **traceoptions** statement at the **[edit protocols rsvp]** hierarchy level. Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory **/var/log**. We recommend that you place RSVP tracing output in the file **rsvp-log**. To examine the contents of the **rsvp-log** file, issue the **file show /var/log/rsvp-log** command.

ResvTear Received Event

LSP Event

ResvTear received

Sample Output

```
user@R1> show mpls lsp extensive
[...Output truncated...]
 23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6
```



```

22 Dec 28 13:46:39 Clear Call
21 Dec 28 13:46:39   ResvTear received
20 Dec 28 13:46:39 Down
19 Dec 28 13:46:39 10.1.13.2: Session preempted
18 Dec 28 13:42:07 Selected as active path
[...Output truncated...]

```

Meaning

This LSP event indicates that an RSVP ResvTear message was received. ResvTear messages remove RSVP reservation states along a path. These messages travel upstream toward senders of the session. This message usually appears in the middle of a run of messages that tear the LSP down.

Cause

In some cases, an ResvTear event is received because a router's reservation state times out. In other cases, when the downstream link fails, the upstream node must eliminate all RSVP states and initiates a ResvTear event. If you are running Fast ReRoute, the upstream node initiates a PathErr message, not a ResvTear message. It is beyond the scope of this document to include all possible reasons for an ResvTear event.

Action

Analyze the status to determine if this is the required behavior. If this is not the required behavior, verify the surrounding LSP events to identify the cause of the problem.

Session Preempted Event

LSP Event

Session preempted

Sample Output 1

```

user@R1> show mpls lsp extensive
[...Output truncated...]
 21 Dec 28 13:46:39 ResvTear received
 20 Dec 28 13:46:39 Down
 19 Dec 28 13:46:39   10.1.13.2: Session preempted
 18 Dec 28 13:42:07 Selected as active path
 17 Dec 28 13:42:07 Record Route:  10.1.13.2 10.1.36.2
[...Output truncated...]

```


Sample Output 2

```

user@R1> show mpls lsp extensive
[...Output truncated...]
 66 Jan  5 09:58:33 CSPF failed: no route toward 10.0.0.6
 65 Jan  5 09:58:33 Clear Call
 64 Jan  5 09:58:33   Session preempted
 63 Jan  5 09:58:33 Down
 62 Jan  5 09:58:32   CSPF failed
    : no route toward 10.0.0.6[2 times]
61 Jan  5 09:57:55 10.1.36.2: Explicit Route: wrong delivery
 60 Jan  5 09:57:34 CSPF failed: no route toward 10.0.0.6[2 times]
 59 Jan  5 09:57:28 CSPF: link down/deleted
10.1.36.1(R3.00/10.0.0.3)->10.1.36.2(R6.00/10.0.0.6)
[...Output truncated...]

```

Meaning

This LSP event indicates that the LSP session was taken over. Sample Output 1 shows the IP address (10.1.13.2) included with the event, indicating the IP address of the router that sent the message. Sample Output 2 does not include an IP address, indicating that the message originated on the ingress router.

Cause

The state of the network might have changed, as shown in Sample Output 1, or an LSP with a higher priority might be using the bandwidth of the LSP.

Action

Refer to the events preceding this event in the history log for more information on what might have caused the preemption. For example, in line 62, the **CSPF failed** message may indicate that you specified a disable constrained-path (**no-cspf**) LSP and an explicit route address that is strict and not directly connected. Additionally, the egress router might have changed its configuration, making the destination address unreachable.

Displaying General LSP Error Events

Purpose

Display extensive information about LSPs, including the 50 most recent history events and the possible reasons why an LSP failed.

Action

To examine error messages, enter the following Junos OS command-line interface (CLI) operational mode command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output

```
user@R1# run show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn, No-decrement-ttl
    Bandwidth: 100Mbps
  14 Jan 21 15:43:39 Requested bandwidth unavailable[3 times]
  13 Jan 21 15:43:21 Deselected as active
  12 Jan 21 15:43:21   Requested bandwidth unavailable
  11 Jan 21 15:43:21 Clear Call
  10 Jan 21 15:42:32 Selected as active path
   9 Jan 21 15:42:32 Record Route:  10.1.12.2 10.1.26.2
   8 Jan 21 15:42:32 Up
[...Output truncated...]
```

Meaning

The sample output from ingress router R1 is a section from the complete output. Typically, the output includes LSP events that led to an LSP failure and the 50 most recent state events. Only one example of a general LSP error event is displayed because it is impossible to provide all of the events described in this topic in one sequence of log history.

For completeness, events not generated by the example network used throughout this book are described to show LSP events that might occur in your network. The output for these events includes the prompt **user@host** rather than the usual **user@R1** prompt.

Admission Control Failure Event

LSP Event

Admission control failure

Sample Output

Not available.

Meaning

This LSP error event indicates that a Resource Reservation Protocol (RSVP) Admission control failure occurred along the LSP path. This event is logged because of an error notification (PathErr message) received from RSVP for the label-switched path.

Cause

This LSP event is caused by inadequate bandwidth on a link along the LSP path. The available bandwidth could not satisfy the requested traffic parameters and no other sessions were pre-empted to accommodate this request.

Action

This error event is not generated by Juniper Networks routers. However, when this event is received by a Juniper Networks router, it appears in the log output of the **show mpls lsp extensive** command.

Explicit Route: Bad Loose Route Event

LSP Event

Explicit Route: bad loose route

Sample Output 1

```
user@R1# run show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
From: 10.0.0.1, State: up, ActiveRoute: 0, LSPname: R1-R6-3
ActivePath: R6-3-1 (secondary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primry R6-3 State: Dn
10 Feb 15 21:21:58 Explicit Route: bad loose route[2 times]
9 Feb 15 21:21:52 Deselected as active
```



```

8 Feb 15 21:21:51 Explicit route: bad loose route
7 Feb 15 21:21:51 10.1.15.1: MPLS label allocation failure
6 Feb 15 21:21:51 MPLS label allocation failure
5 Feb 15 21:21:51 Down
4 Feb 15 21:20:55 Selected as active pathe
3 Feb 15 21:20:55 Record Route: 10.1.15.2 10.1.56.2
2 Feb 15 21:20:55 Up
1 Feb 15 21:20:55 Originate Call
*Secondary R6-3-1 State: Up
Received RRO (ProtectionFlag 1 = Available 2 = InUse 4 = B/W 8 = Node
10 = SoftPreempt):
10.1.12.2 10.1.26.2
4 Feb 15 21:21:52 Selected as active path
3 Feb 15 21:21:52 Record Route: 10.1.12.2 10.1.26.2
2 Feb 15 21:21:52 Up
1 Feb 15 21:21:52 Originate Call
Created: Tue Feb 15 21:20:55 2005
Total 3 displayed, Up 2, Down 1

```

Sample Output 2

```

user@R1# run show protocols mpls
label-switched-path R1-to-R6 {
to 10.0.0.6;
no-cspf;
link-protection;
primary to-R6;
}
label-switched-path R1-to-R6-2 {
to 10.0.0.6;
link-protection;
auto-bandwidth {
adjust-interval 300;
minimum-bandwidth 1;
maximum-bandwidth 1k;
}
}
label-switched-path R1-R6-3 {
to 10.0.0.6;
no-cspf; <--Allows a loose ERO
primary R6-3;
secondary R6-3-1;
}

```



```

}
path to-R6 {
  10.1.15.2 strict;
  10.1.56.2 strict;
}
path R6-3 {
  10.1.15.2 loose; <--Loose ERO
}
path R6-3-1 {
  10.1.12.2;
}
interface fxp0.0 {
  disable;
}
interface all;

```

Meaning

This LSP error event indicates that there is an error in the loose hop specified in the Explicit Route Object (ERO) of a Path message received by a label-switched router (LSR) along the LSP path, indicating an LSP setup failure.

Cause

This LSP error event is caused by control plane unreachability or data plane incompatibility.

Action

Check the LSP configuration at the **[edit protocols mpls]** hierarchy level.

Explicit Route: Bad Strict Route Event

LSP Event

Explicit route: bad strict route

Sample Output 1

```

user@R1> show mpls lsp extensive
[...Output truncated...]
  36 Jan  4 18:04:57 CSPF: link down/deleted 10.1.13.1(R1.00/10.0.0.1)
->10.1.13.2(R3.00/10.0.0.3)
  35 Jan  4 18:04:57 CSPF failed: no route toward 10.0.0.6
  34 Jan  4 18:04:57 Clear Call

```



```

33 Jan  4 18:04:57  Explicit Route: bad strict route
32 Jan  4 18:04:57 No Route toward dest
31 Jan  4 18:04:57 Down
[...Output truncated...]

```

Sample Output 2

```

user@host> show mpls lsp extensive
Ingress LSP: 34 sessions

10.172.2.99
From: 10.172.162.18, State: Up, ActiveRoute: 3726, LSPname:
dcr2.den_to_dcr1.chd_P
ActivePath: Pl_dcr2.den_to_dcr1.chd (primary)
LoadBalance: Random
Metric: 25
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary Pl_dcr2.den_to_dcr1.chd State: Up
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
10=SoftPreempt):
10.70.192.134
16 Jun 28 18:27:51 Selected as active path
15 Jun 28 18:27:51 Record Route: 10.70.192.134
14 Jun 28 18:27:51 Up
13 Jun 28 18:27:29 Deselected as active
12 Jun 28 18:27:28 No Route toward dest
11 Jun 28 18:27:28 Down
10 Jun 18 03:52:18 Selected as active path
9 Jun 18 03:52:18 Record Route: 10.70.192.134
8 Jun 18 03:52:18 Up
7 Jun 18 03:52:18 Originate Call
6 Jun 18 03:52:18 Clear Call
5 Jun 18 03:52:18 Deselected as active
4 Jun 18 02:56:25 Selected as active path
3 Jun 18 02:56:25 Record Route: 10.70.192.134
2 Jun 18 02:56:25 Up
1 Jun 18 02:56:25 Originate Call
Standby Bl_dcr2.den_to_dcr1.chd State: Dn
18 Jun 29 12:49:21 10.70.192.26: Routing loop detected[4798 times]
17 Jun 27 00:53:42 10.70.192.77:
Explicit Route: bad strict route
[20 times]

```



```
16 Jun 27 00:39:49 204.70.192.26: Routing loop detected [3370 times]
[...Output truncated...]
```

Meaning

This LSP event indicates that a poorly formed ERO was generated. Sample Outputs 1 and 2 show that this LSP event was caused by different situations described below.

Cause

This LSP event can be caused by several factors:

- A strict hop address specified for an LSP on a link that does not have RSVP enabled.
- The **no-cspf** statement included in the LSP configuration.
- An error with the configuration of constraints on a Constrained Shortest Path First (CSPF) LSP generates the **CSPF: No route towards dest** message, followed by the **Explicit Route: bad strict route** event.
- An ERO that causes a routing loop. See Sample Output 2.

Action

Examine the strict hop address, remove the **no-cspf** statement, or examine the path and verify that RSVP is enabled on each interface.

Explicit Route: Format Error Event

LSP Event

Explicit route: format error

Sample Output

```
user@R1> show mpls lsp extensive
[...Output truncated...]
10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary    to-R6                State: Dn, No-decrement-ttl
    5 Jan 21 14:37:06 10.1.34.2: Explicit Route: format error [2 times]
    4 Jan 21 14:37:03 Originate Call
```



```
3 Jan 21 14:37:03 Clear Call
[...Output truncated...]
```

Meaning

This LSP event indicates an LSP setup failure in which a Path message error in the the ERO was received by a router along the LSP path.

Cause

This LSP event can be caused by several factors:

- An incorrectly formed ERO in the RSVP Path message.
- A strict hop address specified in the middle of an ERO that is not contiguous.
- An unsupported subobject in the ERO of a router along the LSP path.
- The hop indicated by the RSVP hop object does not match the hop indicated by the ERO.

Action

Examine the strict hop address configuration and make any necessary changes.

Explicit Route: Wrong Delivery Event

LSP Event

Explicit route: wrong delivery

Sample Output 1

```
user@host> show mpls lsp extensive
[...Output truncated...]
Primary use-TOKYO State: Dn, No-decrement-ttl
  3 Sep 19 00:25:45 10.222.45.2: Explicit Route: wrong delivery
  2 Sep 19 00:25:34 No Route[8 times]
  1 Sep 19 00:23:01 Originate Call
[...Output truncated...]
```

Sample Output 2

```
user@host> show mpls lsp extensive
[...Output truncated...]
```



```

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary    to-R6                State: Dn
    40 Jan 26 16:35:26  10.1.36.2: Explicit Route: wrong delivery [2 times]
    39 Jan 26 16:35:23 Originate Call
    38 Jan 26 16:35:23 Clear Call
[...Output truncated...]

```

Sample Output 3

```

user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
  to 10.0.0.6;
  no-cspf;
  primary to-R6;
}
path to-R6 {
  10.1.13.2 strict;
  10.1.56.1 strict;      <<< IP address not directly connected to 10.1.13.2
  10.1.26.1 strict;
}

```

Meaning

This LSP event indicates that a RSVP message with an ERO arrived at the wrong router, even though a strict route was specified. The receiving router determines that the address is inconsistent with the ERO, and generates the error message. Note that the IP address of the sending router precedes the error event; for example, **10.222.45.2** in Sample Output 1, and **10.1.36.2** in Sample Output 2.

Cause

This LSP event can be caused by several factors:

- The loopback (**lo0**) interface on the ingress router is not configured at the **[edit protocols isis]** hierarchy level. After the loopback (**lo0**) interface is included in the Intermediate System-to-Intermediate System (IS-IS) configuration, and while IS-IS is forming adjacencies, an RSVP packet is forwarded to an incorrect destination, **10.222.45.2**, as shown in Sample Output 1.
- A strict path is configured to a directly connected router, then another strict path is configured to an IP address that is not directly connected. For example, Sample Output 3 shows that the path **to-R6** includes three IP addresses, one of which (**10.1.56.1**) is not directly connected to the other IP addresses in the path.

Action

Take appropriate action. On the ingress router, include the loopback (lo0) interface at the [edit protocols isis] hierarchy level, change the definition of the strict path at the [edit protocols mpls path *path-name*] hierarchy level, or verify the validity of all IP addresses listed in the named path referenced by the LSP hop by hop.

Invalid Destination Address Event

LSP Event

Invalid Dest addr

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Metric: 100
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    4 Apr 22 10:22:15   Invalid Dest Addr
    3 Apr 22 10:22:15   Originate Call
    2 Apr 22 10:22:15   Invalid Dest Addr
    1 Apr 22 10:22:15   Originate Call
  Created: Fri Apr 22 10:22:16 2005
Total 1 displayed, Up 0, Down 1
```

Meaning

This LSP event indicates that the **to** address configured at the [edit protocols mpls labeled-switched-path *name*] hierarchy level is invalid.

Cause

This LSP event is caused when the **to** address of the LSP is the loopback address of the ingress router. A contributing factor may be that the **no-cspf** statement is included in the LSP configuration.

Action

Verify that the LSP destination address is not the local router's loopback address, and check that the addresses on the local router are correctly configured.

Invalid Filter for Policing Event

LSP Event

Invalid filter for policing

Sample Output

Not available. This LSP event indicates an abnormal condition and is difficult to recreate.

Meaning

Although a policer was configured on the LSP, the corresponding firewall filter index was not found, indicating a failure in the routing protocol process (rpd) or the firewall process (dfwd).

Cause

A possible cause is that the routing protocol process (rpd) or the firewall process (dfwd) were restarted in a situation in which the LSP was established.

Action

Not applicable.

MPLS Graceful Restart: Recovery Failed Event

LSP Event

MPLS graceful restart: recovery failed

Sample Output

Not available.

Meaning

This LSP event indicates unsuccessful recovery of an LSP path after graceful restart, resulting in potential traffic loss.

Cause

This LSP event is caused by several factors:

- MPLS graceful restart procedures may have been aborted by this LSR.
- MPLS graceful restart is disabled, by configuration, during the recovery period.
- An MPLS LSP path is disabled either due to a configuration change or due to an error during the recovery period.
- CSPF computation failed for the restarted LSP path with parameters and constraints preserved across the restart.

- A signaling failure occurred and an RSVP PathErr was received on the LSP path signaled after a restart.
- A network failure occurred on some hop that the LSP was traversing during the recovery period.

Action

Check the MPLS logs for more details about the failure.

MPLS Label Allocation Failure Event

LSP Event

MPLS label allocation failure

Sample Output

```
user@R1> show mpls lsp extensive
[...Output truncated...]
 24 Jan 20 09:25:35 CSPF failed: no route toward 10.0.0.6
 23 Jan 20 09:25:35 Clear Call
 22 Jan 20 09:25:35 Deselected as active
 21 Jan 20 09:25:35 10.1.13.1: MPLS label allocation failure
 20 Jan 20 09:25:34 MPLS label allocation failure
 19 Jan 20 09:25:34 Down
[...Output truncated...]
```

Meaning

This LSP event indicates that the MPLS protocol or the **family mpls** statement were not configured properly. When the LSP event is preceded by an IP address, the address is typically the router that has the MPLS configuration error.

Cause

This LSP event is caused by the omission of interfaces at the **[edit protocols mpls]** hierarchy level or failure to configure the **family mpls** statement at the **[edit interfaces type-fpc/pic/port]** hierarchy level. The **family mpls** statement specifies to the interface ASICs to permit protocol code 0x8847 (unicast MPLS) into the router.

Action

Include interfaces at the **[edit protocols mpls]** hierarchy level, or include the **family mpls** statement at the **[edit interfaces type-fpc/pic/port]** hierarchy level. You must configure the **family mpls** statement, in the same way that you must configure the **family iso** statement for IS-IS.

NOTE: Do not configure the **family mpls** statement on the loopback (lo0) interface.

Non-RSVP Capable Router Detected Event

LSP Event

Non-RSVP capable router detected

Sample Output

```
user@host> show mpls lsp extensive
[...Output truncated...]
10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn, No-decrement-ttl
    19 Jan 21 15:05:37 10.1.24.2:  Non-RSVP capable router detected
    18 Jan 21 15:04:52 10.1.26.2:  Non-RSVP capable router detected [4 times]
    17 Jan 21 15:04:34 Originate Call
    16 Jan 21 15:04:34 Clear Call
[...Output truncated...]
```

Meaning

This LSP event indicates that a router, forwarding packets to the egress router, was not configured for RSVP.

Cause

This LSP event might be caused when a router not configured for RSVP forwards an RSVP packet toward the egress router without decrementing the **Send_TTL** value in the RSVP common header. The next downstream router detects that the **Send_TTL** value and the **IP_TTL** value are different, and generates this LSP event. Note that two different routers generated the same error message at different times.

Action

Configure the router in question with RSVP.

No Route Toward Destination Event

LSP Event

No route toward destination

Sample Output 1

```
user@R1> show mpls lsp extensive
[...Output truncated...]
 35 Oct 26 22:48:36 Down
 34 Oct 26 22:48:29 CSPF failed: no route toward 10.0.0.1[4 times]
 33 Oct 26 22:47:25 CSPF: link down/deleted
10.1.13.2(R3.00/10.0.0.3)->10.1.13.1(R1.00/10.0.0.1)
 32 Oct 26 22:47:25 CSPF failed: no route toward 10.0.0.1
 31 Oct 26 22:47:25 10.1.36.1: No Route toward dest
 30 Oct 26 22:33:54 Selected as active path
 29 Oct 26 22:33:53 Record Route: 10.1.36.1 10.1.13.1
[...Output truncated...]
```

Sample Output 2

```
user@R1> show mpls lsp extensive
[...Output truncated...]
10.0.0.6
  From: 10.0.0.1 , State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
ActivePath: (none)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary                               State: Dn
  Will be enqueued for recomputation in 7 second(s).
 13 Oct 25 16:29:28 Deselected as active
 12 Oct 25 16:29:27 CSPF failed: no route toward 10.0.0.6
 11 Oct 25 16:29:27 CSPF: link down/deleted
10.1.13.1(R1.00/10.0.0.1)->10.1.13.2(R3.00/10.0.0.3)
 10 Oct 25 16:29:27 CSPF failed: no route toward 10.0.0.6
  9 Oct 25 16:29:27 Clear Call
  8 Oct 25 16:29:27 Explicit Route: bad strict route
  7 Oct 25 16:29:27 No Route toward dest
  6 Oct 25 16:29:27 Down
[...Output truncated...]
```


Meaning

This LSP event indicates that the router at address **10.1.36.1** in Sample Output 1 does not have a route to the specified destination. Sample Output 2 shows that the local router, ingress router **10.0.0.1**, does not have a route to the specified destination.

Cause

This LSP event is caused by different factors. The egress interface of a transit router might not have RSVP enabled, or IP reachability to the destination (either the egress router or the next address in the ERO) does not exist.

Action

Enable RSVP on the transit router's egress interface, or examine the IP configuration of the relevant router.

Unsupported Traffic Class Event

LSP Event

Unsupported traffic class

Sample Output

Not available. This LSP event indicates an abnormal condition and is difficult to recreate.

Meaning

This LSP error event is a Juniper Networks proprietary error, indicating that a DiffServ-traffic engineering (TE) LSP was signaled with one or more traffic classes with values greater than the four traffic classes currently supported.

Cause

Not available.

Action

Not available.

CSPF: Computation Result Accepted Event

LSP Event

CSPF: computation result accepted

Sample Output

```
user@R1> show mpls lsp extensive
[...Output truncated...]
 57 Jan  5 09:54:37 Record Route:  10.1.13.2 10.1.36.2
 56 Jan  5 09:54:37 Up
 55 Jan  5 09:54:37 Originate Call
 54 Jan  5 09:54:37  CSPF: computation result accepted
[...Output truncated...]
```

Meaning

This LSP event indicates that CSPF pruned the traffic engineering database of noncompliant links and found a shortest path. CSPF generated an ERO, which was then passed to the RSVP.

Cause

Not applicable.

Action

Not applicable

CSPF: Reroute Due to Re-Optimization Event

LSP Event

CSPF: Reroute due to re-optimization

Sample Output

```
user@host> show mpls lsp extensive
[...Output truncated...]
 9 Dec 11 17:32:35  Up
   8 Dec 11 17:32:35  Clear Call
   7 Dec 11 17:32:35  CSPF: computation result accepted
   6 Dec 11 17:32:35  CSPF: Reroute due to re-optimization
   5 Dec 11 17:28:29  CSPF: computation result ignored
   4 Dec 11 17:24:23  Record Route:  10.35.38.2 S 192.168.135.29 S
10.35.39.1 S 10.35.40.2 S 10.35.41.1 S
   3 Dec 11 17:24:23  Up
[...Output truncated...]
```


Meaning

This LSP event indicates that CSPF found an optimal path for LSP traffic, and switched over to the new path.

Cause

This is a periodic or one-time reoptimization event.

Action

Not applicable.

Retry Limit Exceeded Event

LSP Event

Retry limit exceeded

Sample Output 1

```
user@R1> show mpls lsp extensive
[...Output truncated...]
10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    RetryCount: 13
    RetryLimit: 1
    12 Jan 14 15:39:30 Clear Call
    11 Jan 14 15:39:30 Retry limit exceeded
    10 Jan 14 15:39:10 10.1.12.1: MPLS label allocation failure[11 times]
[...Output truncated...]
```

Sample Output 2

```
user@R1> show mpls lsp extensive
[...Output truncated...]
10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
```



```

Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary                               State: Dn
14 Jan 14 15:44:07 10.1.12.1: MPLS label allocation failure[3 times]
13 Jan 14 15:43:58 Originate Call
12 Jan 14 15:39:30 Clear Call
11 Jan 14 15:39:30   Retry limit exceeded
10 Jan 14 15:39:10 10.1.12.1: MPLS label allocation failure[11 times]
[...Output truncated...]

```

Meaning

This LSP event indicates that the number of CSPF path computations for a particular path exceeded a configured retry limit. After this point, the path is not recomputed or signaled, unless the user intervenes.

Cause

The number of CSPF path computations for an LSP path exceeded the configured non-zero retry limit. Sample Output 1 shows that a configured retry limit of 1 was exceeded by the retry count of 13.

Action

Enter the **clear mpls lsp** command to disconnect and restart the LSP. Sample Output 2 shows that events 13 and 14 were generated after the **clear mpls lsp** command was issued. This operation disconnects existing RSVP sessions on the ingress router, releases the routes and states associated with the LSPs, and starts a new LSP. Issuing this command might impact traffic travelling along the LSP, because of a time lag that can occur between the old path being torn down and the new path being set up.

Log the Software Version Information

Purpose

The purpose of this action is to log the Junos OS version information.

Action

Use the following Junos OS CLI operational mode command:

```
user@host> show version | save filename
```

Sample Output

```

user@host> show version | save test
Wrote 39 lines of output to 'test'

```



```

user@host> show version
Hostname:  my-router.net
Model:  m10
JUNOS Base OS boot [5.0R5]
JUNOS Base OS Software Suite [5.0R5]
JUNOS Kernel Software Suite [5.0R5]
JUNOS Routing Software Suite [5.0R5]
JUNOS Packet Forwarding Engine Support [5.0R5]
JUNOS Crypto Software Suite [5.0R5]
JUNOS Online Documentation [5.0R5]
KERNEL 5.0R5 #0 built by builder on 2002-03-02 05:10:28 UTC
MGD release 5.0R5 built by builder on 2002-03-02 04:45:32 UTC
CLI release 5.0R5 built by builder on 2002-03-02 04:44:22 UTC
CHASSISD release 5.0R5 built by builder on 2002-03-02 04:43:37 UTC
DCD release 5.0R5 built by builder on 2002-03-02 04:42:47 UTC
RPD release 5.0R5 built by builder on 2002-03-02 04:46:17 UTC
SNMPD release 5.0R5 built by builder on 2002-03-02 04:52:26 UTC
MIB2D release 5.0R5 built by builder on 2002-03-02 04:45:37 UTC
APSD release 5.0R5 built by builder on 2002-03-02 04:43:31 UTC
VRRPD release 5.0R5 built by builder on 2002-03-02 04:52:34 UTC
ALARMD release 5.0R5 built by builder on 2002-03-02 04:43:24 UTC
PFED release 5.0R5 built by builder on 2002-03-02 04:46:06 UTC
CRAFTD release 5.0R5 built by builder on 2002-03-02 04:44:30 UTC
SAMPLED release 5.0R5 built by builder on 2002-03-02 04:52:20 UTC
ILMID release 5.0R5 built by builder on 2002-03-02 04:45:21 UTC
BPRELAYD release 5.0R5 built by builder on 2002-03-02 04:42:41 UTC
RMOPD release 5.0R5 built by builder on 2002-03-02 04:46:11 UTC
jkernel-dd release 5.0R5 built by builder on 2002-03-02 04:41:07 UTC
jroute-dd release 5.0R5 built by builder on 2002-03-02 04:41:21 UTC
jdocs-dd release 5.0R5 built by builder on 2002-03-02 04:39:11 UTC

```

Meaning

The sample output shows the hostname, router model, and the different Junos OS packages, processes, and documents.

Log the Hardware Version Information

Purpose

You should log hardware version information in the rare event that a router cannot successfully reboot and you cannot obtain the Routing Engine serial number. The Routing Engine serial number is necessary for Juniper Networks Technical Assistance Center (JTAC) to issue a return to manufacturing authorization

(RMA). Without the Routing Engine serial number, an onsite technician must be dispatched to issue the RMA.

Action

To log the router chassis hardware version information, use the following Junos OS CLI operational mode command:

```
user@host> show chassis hardware | save filename
```

Sample Output

The output for the M-series routers varies depending on the chassis components of each router. All routers have a chassis, midplanes or backplanes, power supplies, and Flexible PIC Concentrators (FPCs). Refer to the hardware guides for information about the different chassis components.

```
user@host> show chassis hardware | save test
Wrote 43 lines of output to 'test'
```



```
user@host> show chassis hardware
```

Item	Version	Part number	Serial number	Description
Chassis			101	M160
Midplane	REV 02	710-001245	S/N AB4107	
FPM CMB	REV 01	710-001642	S/N AA2911	
FPM Display	REV 01	710-001647	S/N AA2999	
CIP	REV 02	710-001593	S/N AA9563	
PEM 0	Rev 01	740-001243	S/N KJ35769	DC
PEM 1	Rev 01	740-001243	S/N KJ35765	DC
PCG 0	REV 01	710-001568	S/N AA9794	
PCG 1	REV 01	710-001568	S/N AA9804	
Host 1			da000004f8d57001	teknor
MCS 1	REV 03	710-001226	S/N AA9777	
SFM 0 SPP	REV 04	710-001228	S/N AA2975	
SFM 0 SPR	REV 02	710-001224	S/N AA9838	Internet Processor I
SFM 1 SPP	REV 04	710-001228	S/N AA2860	
SFM 1 SPR	REV 01	710-001224	S/N AB0139	Internet Processor I
FPC 0	REV 03	710-001255	S/N AA9806	FPC Type 1
CPU	REV 02	710-001217	S/N AA9590	
PIC 1	REV 05	750-000616	S/N AA1527	1x OC-12 ATM, MM
PIC 2	REV 05	750-000616	S/N AA1535	1x OC-12 ATM, MM
PIC 3	REV 01	750-000616	S/N AA1519	1x OC-12 ATM, MM
FPC 1	REV 02	710-001611	S/N AA9523	FPC Type 2
CPU	REV 02	710-001217	S/N AA9571	
PIC 0	REV 03	750-001900	S/N AA9626	1x STM-16 SDH, SMIR
PIC 1	REV 01	710-002381	S/N AD3633	2x G/E, 1000 BASE-SX

FPC 2				FPC Type OC192
CPU	REV 03	710-001217	S/N AB3329	
PIC 0	REV 01			1x OC-192 SM SR-2

Meaning

The sample output shows the hardware inventory for an M160 router with a chassis serial number of 101. For each component, the output shows the version number, part number, serial number, and description.

Log the System Boot-Message Information

Action

To log the system boot-message information, use the following Junos OS CLI operational mode command:

```
user@host> show system boot-messages | save filename
```

Sample Output

```
user@host> show system boot-messages | save test
Wrote 80 lines of output to 'test'

user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
    The Regents of the University of California. All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC

tlim@device1.example.com:/p/build/20000216-0905/4.1/release_kernel/sys/compile/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
    Origin = "GenuineIntel"  Id = 0x66a  Stepping=10

Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 60000
```



```

on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on
pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 60400
on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on
pci0:19:0
Probing for devices on PCI bus 1:mcs0 <Miscellaneous Control Subsystem> rev 12
class ff0000 int a irq 12 on pci1:13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on
pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300
ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctlr(0) TI PCI-1131 [CardBus bridge model] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SDCFB-80>, single-sector-i/o
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>

```



```

wdl: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300
fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wdls1b as swap device
Automatic reboot in progress...
/dev/rwd0sla: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0sle: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0sla: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwdls1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)

```

Meaning

The sample output shows the initial messages generated by the system kernel upon boot. This is the content of the `/var/run/dmesg.boot` file.

Log the BGP, IS-IS, and OSPF Adjacency Information

Purpose

The following commands log useful information about Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Open Shortest Path First (OSPF) protocols. If you have other protocols installed, such as Multiprotocol Label Switching (MPLS), Resource Reservation Protocol (RSVP), or Protocol Independent Multicast (PIM), you also might log summary information for them.

Action

To log the protocol peer information, use the following Junos OS CLI operational mode commands:

```

user@host> show bgp summary | save filename
user@host> show isis adjacency brief | save filename
user@host> show ospf neighbor brief | save filename

```

Sample Output 1

```

user@host> show bgp summary | save test
Wrote 45 lines of output to 'test'

user@host> show bgp summary

```



```

Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                4          4          0          0          0          0
Peer            AS        InPkt      OutPkt    OutQ    Flaps Last Up/Dwn
State|#Active/Received/Damped..
9.9.3.1          2        2627      2628      0        0    21:50:12 4/4/0
                0/0/0

```

Sample Output 2

```

user@host> show isis adjacency brief | save test
Wrote 7 lines of output to 'test'

```

```

user@host> show isis adjacency brief
IS-IS adjacency database:
Interface  System          L State          Hold (secs) SNPA
so-1/0/0.0 1921.6800.5067 2 Up             13
so-1/1/0.0 1921.6800.5067 2 Up             25
so-1/2/0.0 1921.6800.5067 2 Up             20
so-1/3/0.0 1921.6800.5067 2 Up             19
so-2/0/0.0 1921.6800.5066 2 Up             19
so-2/1/0.0 1921.6800.5066 2 Up             17
so-2/2/0.0 1921.6800.5066 2 Up             20
so-2/3/0.0 1921.6800.5066 2 Up             20
so-5/0/0.0 ranier          2 Up             17

```

Sample Output 3

```

user@host> show ospf neighbor brief | save test
Wrote 10 lines of output to 'test'

```

```

user@host> show ospf neighbor brief
Address      Intf          State      ID           Pri  Dead
10.168.254.225 fxp3.0       2Way       10.250.240.32 128  36
10.168.254.230 fxp3.0       Full       10.250.240.8  128  38
10.168.254.229 fxp3.0       Full       10.250.240.35 128  33
10.1.1.129      fxp2.0       Full       10.250.240.12 128  37
10.1.1.131      fxp2.0       Full       10.250.240.11 128  38
10.1.2.1        fxp1.0       Full       10.250.240.9  128  32
10.1.2.81       fxp0.0       Full       10.250.240.10 128  33

```


Meaning

Sample output 1 displays summary information about BGP and its neighbors. Sample output 2 displays information about IS-IS neighbors. Sample output 3 displays information about all OSPF neighbors.

Back Up the Currently Running and Active File System

To back up the currently running and active file system so that you can recover to a known, stable environment in case there is a problem during the reinstall, use the following operational mode command:

```
user@host> request system snapshot
```

Sample Output

```
user@host> request system snapshot
```

```
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)
umount: /altconfig: not currently mounted
Copying /config to /altconfig.. (this may take a few minutes)
The following filesystems were archived: / /config
```

Meaning

On devices running Junos OS, this command backs up the root file system to **/altroot**, and **/config** is backed up to **/altconfig**. The root and **/config** file systems are on the router's internal flash drive, and the **/altroot** and **/altconfig** file systems are on the router's hard drive.

On devices running Junos OS Evolved, this command only copies the contents of the **/soft** directory, which includes the running version of Junos OS Evolved.

NOTE: After you issue the **request system snapshot** command, you cannot return to the previous version of the software because the running and backup copies of the software are identical.

Reinstall Junos OS

Action

To reinstall Junos OS, follow these steps:

1. Insert the removable medium (boot floppy) into the router.
2. Reboot the router, either by power-cycling it or by issuing the **request system reboot** command from the CLI.
3. At the following prompt, type **y**:

```
WARNING: The installation will erase the contents of your disk. Do you wish to  
continue (y/n)?
```

The router copies the software from the removable medium onto your system, occasionally displaying status messages. This can take up to 10 minutes.

4. Remove the removable medium when prompted.

The router reboots from the primary boot device on which the software is installed. When the reboot is complete, the router displays the login prompt.

Reconfigure Junos OS

After you have reinstalled the software, you must copy the router's configuration files back to the router. (You also can configure the router from scratch, as described in *Junos System Basics Configuration Guide*) However, before you can copy the configuration files, you must establish network connectivity.

To reconfigure the software, follow these steps:

1. [Configure Host Names, Domain Names, and IP Addresses | 2356](#)
2. [Protecting Network Security by Configuring the Root Password | 2357](#)
3. [Check Network Connectivity | 2359](#)

Configure Host Names, Domain Names, and IP Addresses

To configure the machine name, domain name, and various addresses, follow these steps:

1. Log in as **root**. There is no password.
2. Start the CLI:

```
root# cli  
root@>
```


3. Enter configuration mode:

```
cli> configure  
[edit]  
root@#
```

4. Configure the name of the machine. If the name includes spaces, enclose the entire name in quotation marks (" "):

```
[edit]  
root@# set system host-name host-name
```

5. Configure the machine's domain name:

```
[edit]  
root@# set system domain-name domain-name
```

6. Configure the IP address and prefix length for the router's management Ethernet interface:

```
[edit]  
root@# set interfaces fxp0 unit 0 family inet address address / prefix-length
```

7. Configure the IP address of a default router. This system is called the backup router because it is used only while the routing protocol process is not running.

```
[edit]  
root@# set system backup-router address
```

8. Configure the IP address of a Domain Name Server (DNS) server:

```
[edit]  
root@# set system name-server address
```

Protecting Network Security by Configuring the Root Password

Configuring the root password on your Junos OS-enabled router helps prevent unauthorized users from making changes to your network. The root user (also referred to as superuser) has unrestricted access and full permissions within the system, so it is crucial to protect this account by setting a strong password when setting up a new router.

After a new router is initially powered on, you log in as the user **root** with no password. Junos OS requires configuration of the root password before it accepts a commit operation.

To set the root password, you have a few options as shown in the following procedure.

- Enter a plain-text password that Junos OS encrypts.
- Enter a password that is already encrypted.
- Enter a secure shell (ssh) public key string.

The most secure options of these three are using an already encrypted password or an ssh public key string. Pre-encrypting your password or using a ssh public key string means the plain-text version of your password will never be transferred over the internet, protecting it from being intercepted by a man-in-the-middle attack.

BEST PRACTICE: Optionally, instead of configuring the root password at the **[edit system]** hierarchy level, you can use a configuration group to strengthen security.

To set the root password:

1. Use one of these methods to configure the root password:

- To enter a plain-text password that the system encrypts for you:

```
[edit groups global system]
root@# set root-authentication plain-text-password
New Password: type password here
Retype new password: retype password here
```

As you enter a plain-text password into the CLI, Junos OS hides it from view and encrypts it immediately. You do not have to configure Junos OS to encrypt the password as in some other systems. In the resulting configuration, the encrypted password is marked as **## SECRET-DATA** so that it cannot be seen.

- To enter a password that is already encrypted:



CAUTION: Do not use the **encrypted-password** option unless the password is *already* encrypted, and you are entering the encrypted version of the password.

If you accidentally configure the **encrypted-password** option with a plain-text password or with blank quotation marks (" "), you will not be able to log in to the device as root, and you will need to complete the root password recovery process.


```
[edit groups global system]
root@# set root-authentication encrypted-password password
```

- To enter an ssh public key string:

```
[edit groups global system]
root@# set root-authentication (ssh-dsa | ssh-eccdsa | ssh-rsa key)
```

2. If you used a configuration group, apply it with the command **set apply-groups**, replacing *<group name>* with the configuration group name.

For example:

```
[edit]
user@host# set apply-groups <group name>
```

3. Commit the changes.

```
root@# commit
```

SEE ALSO

| *Accessing a Junos OS Device the First Time*

Check Network Connectivity

Purpose

Establish that the router has network connectivity.

Action

To check that the router has network connectivity, issue a **ping** command to a system on the network:

```
root@> ping address
```

If there is no response, verify that there is a route to the **address** using the **show route** command. If the address is outside your **fxp0** subnet, add a static route. Once the backup configuration is loaded and committed, the static route is no longer needed and should be deleted.

Configure Host Names, Domain Names, and IP Addresses

To configure the machine name, domain name, and various addresses, follow these steps:

1. Log in as **root**. There is no password.
2. Start the CLI:

```
root# cli
root@>
```

3. Enter configuration mode:

```
cli> configure
[edit]
root@#
```

4. Configure the name of the machine. If the name includes spaces, enclose the entire name in quotation marks (" "):

```
[edit]
root@# set system host-name host-name
```

5. Configure the machine's domain name:

```
[edit]
root@# set system domain-name domain-name
```

6. Configure the IP address and prefix length for the router's management Ethernet interface:

```
[edit]
root@# set interfaces fxp0 unit 0 family inet address address / prefix-length
```

7. Configure the IP address of a default router. This system is called the backup router because it is used only while the routing protocol process is not running.

```
[edit]
root@# set system backup-router address
```

8. Configure the IP address of a Domain Name Server (DNS) server:

```
[edit]
root@# set system name-server address
```


Check Network Connectivity

Purpose

Establish that the router has network connectivity.

Action

To check that the router has network connectivity, issue a **ping** command to a system on the network:

```
root@> ping address
```

If there is no response, verify that there is a route to the ***address*** using the **show route** command. If the address is outside your **fxp0** subnet, add a static route. Once the backup configuration is loaded and committed, the static route is no longer needed and should be deleted.

Automatic Autobandwidth Adjustment Failed Event

LSP Event

Autobw adjustment failed

Sample Output 1

```
user@R1> show configuration protocols mpls
statistics {
    file auto-bw.log;
    interval 5;
    auto-bandwidth;
}
label-switched-path R1-to-R6 {
    to 10.0.0.6;
    auto-bandwidth {
        adjust-interval 300;
        adjust-threshold 10;
        minimum-bandwidth 5m;
        maximum-bandwidth 80m;
    }
}
label-switched-path R1-to-R3 {
    to 10.0.0.3;
    auto-bandwidth {
        adjust-interval 300;
        adjust-threshold 10;
```



```

        minimum-bandwidth 155m;
        maximum-bandwidth 155m;
    }
}

```

Sample Output 2

```

user@R1> show mpls lsp extensive
Ingress LSP: 3 sessions

10.0.0.3
  From: 10.0.0.1, State: Up, ActiveRoute: 5, LSPname: R1-to-R3
  ActivePath: (primary)
  LoadBalance: Random
  Metric: 1
  Autobandwidth
  MinBW: 155Mbps MaxBW: 155Mbps
  AdjustTimer: 300 secs AdjustThreshold: 10%
  Max AvgBW util: 192bps, Bandwidth Adjustment in 219 second(s).
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 10)
10.1.13.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.13.2
      7 Feb 17 15:41:12 Autobw adjustment failed
      6 Feb 17 15:41:12 CSPF failed: no route toward 10.0.0.3
      5 Feb 17 15:36:23 Selected as active path
      4 Feb 17 15:36:23 Record Route: 10.1.13.2
      3 Feb 17 15:36:23 Up
      2 Feb 17 15:36:23 Originate Call
      1 Feb 17 15:36:23 CSPF: computation result accepted
  Created: Thu Feb 17 15:36:23 2005
[...Output truncated...]

```

Meaning

This LSP event indicates that a periodic (timer-based) autobandwidth adjustment for the LSP is triggered at the end of the adjust interval. The adjustment fails, and the LSP stays up on the existing path with its current bandwidth.

Cause

Adjustment failure may be due to a path CSPF computation failure with the adjust bandwidth or a signaling failure on the new path.

At the end of the time interval specified at the `[edit protocols mpls label-switched-path auto-bandwidth]` hierarchy level, the current maximum average bandwidth usage is compared to 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.

Action

Take action appropriate to the situation:

- Verify the MPLS and RSVP configuration on all available paths to the LSP endpoint.
- Check available bandwidth on alternate paths using the **show rsvp interface** command. If not enough bandwidth is available on any available paths, adjust the minimum-bandwidth parameter for the LSP in order to establish or adjust the priority to allow the LSP to preempt another LSP of lesser priority. For an LSP to be preempted, its hold priority must be lower than the LSP you are trying to establish.

Configuring Automatic Bandwidth Allocation for LSPs

IN THIS SECTION

- [Configuring Automatic Bandwidth Allocation on LSPs | 2364](#)
- [Requesting Automatic Bandwidth Allocation Adjustment | 2370](#)

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

At the end of the automatic bandwidth allocation time interval, the current maximum average bandwidth usage is compared with the allocated bandwidth for the LSP. If the LSP needs more bandwidth, an attempt is made to set up a new path where bandwidth is equal to the current maximum average usage. If the attempt is successful, the LSP's traffic is routed through the new path and the old path is removed. If the attempt fails, the LSP continues to use its current path.

NOTE: In calculating the value for **Max AvgBW** (relative to the ingress LSP), the sample collected during make before break (MBB) is ignored to prevent inaccurate results. The first sample after a bandwidth adjustment, or after a change in the LSP ID (regardless of path change), is also ignored.

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:

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

IN THIS SECTION

- [Configuring the Automatic Bandwidth Allocation Interval | 2365](#)
- [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth | 2366](#)
- [Configuring the Automatic Bandwidth Adjustment Threshold | 2367](#)
- [Configuring a Limit on Bandwidth Overflow and Underflow Samples | 2368](#)
- [Configuring Passive Bandwidth Utilization Monitoring | 2370](#)

To enable automatic bandwidth allocation on an LSP, include the **auto-bandwidth** statement:

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

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 556](#).

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.

NOTE: For a label-switched path (LSP) that has both **bandwidth** and **minimum-bandwidth** for autobandwidth configured under the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level, the LSP bandwidth is adjusted differently.

The LSP is initiated with the bandwidth value configured under the **bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level. At the expiry of the **adjust-interval** timer, the LSP bandwidth gets adjusted based on the traffic flow.

If the bandwidth to be signaled is less than the value configured under the **minimum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the minimum bandwidth.

If the bandwidth to be signaled is greater than the value configured under the **maximum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the maximum bandwidth.

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 resignaled 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. Starting in Junos OS Release 14.1R9, 15.1R7, 16.1R5, 16.1X2, 16.2R3, and 17.2R2, all zero value bandwidth samples are considered as underflow samples, except for the zero value samples that arrive after an LSP comes up for the first time, and the zero value samples that arrive first after a Routing Engine switchover.

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 554](#)).
- 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 554](#)), 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.

Release History Table

Release	Description
14.1R9	Starting in Junos OS Release 14.1R9, 15.1R7, 16.1R5, 16.1X2, 16.2R3, and 17.2R2, all zero value bandwidth samples are considered as underflow samples, except for the zero value samples that arrive after an LSP comes up for the first time, and the zero value samples that arrive first after a Routing Engine switchover.

RELATED DOCUMENTATION

[Configuring MPLS to Gather Statistics | 242](#)

[Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs | 556](#)

[request mpls lsp adjust-autobandwidth | 3142](#)

[show mpls lsp | 3214](#)

[show mpls lsp autobandwidth | 3243](#)

Displaying DiffServ-Aware Traffic-Engineered LSP Events

Purpose

A DiffServ-aware traffic-engineered LSP is configured with a bandwidth reservation for a specific class type, and carries traffic for a single class type. On the packets, the class type is specified by the experimental (EXP) bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, instead of being signaled in Resource Reservation Protocol (RSVP).

The class type must be configured consistently across the DiffServ domain, and must be consistent from router to router in the network. You can unambiguously map a class type to a queue. On each node router, the class-of-service queue configuration for an interface translates to the available bandwidth for a particular class type on that link. For more information about forwarding classes and class of service, see the Junos Class of Service Configuration Guide. For more information about differentiated services, see RFC 3270, *MPLS Applications User Guide Support of Differentiated Services*.

When the configuration of a DiffServ-aware traffic-engineered LSP is incorrect, an even or error message might occur in the output of the **show mpls lsp extensive** command.

Action

To display LSP events that can occur with a DiffServ-aware LSP, enter the following Junos OS command-line interface (CLI) operational mode command from the ingress router:

```
user@host> show mpls lsp extensive
```

Sample Output

Not available.

Unsupported Traffic Class Event

LSP Event

Unsupported traffic class

Sample Output

Not available.

Meaning

This LSP error event is a Juniper Networks proprietary error indicating that a Diffserv traffic engineering LSP was signaled with one or more traffic classes with values greater than the four traffic classes currently supported by the Junos OS.

Cause

Not available.

Action

Not available.

Traffic Class Value Out of Allowed Range Event

LSP Event

Traffic class value out of allowed range

Sample Output

Not available.

Meaning

This LSP error event is a Juniper Networks proprietary error indicating that a single class, IETF-style DiffServ traffic engineering LSP was signaled with a traffic class value of zero, which is invalid.

Cause

Not available.

Action

Not available.

The Combination of Setup Priority and Traffic Class Is Not One of the Configured TE Classes Event

LSP Event

The combination of setup-priority and traffic class is not one of the configured TE-classes

Sample Output

Not available.

Meaning

This LSP error event is a Juniper Networks proprietary error that indicates the setup priority signaled in the Path message for the LSP does not match the supported DiffServ traffic engineering classes configured on a label-switching router (LSR) along the LSP path.

Cause

This LSP error event is caused by incorrect configuration of the LSP setup priority on the ingress LSR, or the incorrect configuration of a DiffServ traffic engineering class on an LSR along the LSP path.

Action

Correct the configuration depending on the supported traffic engineering classes.

RSVP Error, Subcode 7, Signal Type Does Not Match Link Encoding Event

LSP Event

RSVP error, subcode 7, signal-type does not match link encoding

Sample Output

Not available

Meaning

This LSP error event is a Juniper Networks proprietary error reported for GMPLS LSPs when the configured signal bandwidth does not match the encoding type of the traffic engineering link selected on the first hop.

Cause

An incorrect Sender Tspec is used with a particular LSP switching or encoding type.

Action

Not available.

Unacceptable Label Value Event

LSP Event

Unacceptable label value event

Sample Output

Not available.

Meaning

This LSP error event indicates that the label value signaled in either the Path or Resv message was unacceptable to a label-switched router (LSR) along the LSP path.

Cause

For GMPLS LSPs, this LSP error event is generated by incorrect label mapping configured on one of the LSRs, or by deletion of a resource that was being used by an LSP.

Action

Not available.

Unsupported Switching Type Event

LSP Event

Unsupported switching type

Sample Output

Not available.

Meaning

This LSP error event indicates that the switching type requested in the generalized label request for a GMPLS LSP is unsupported on the corresponding selected traffic engineering link.

Cause

Not available.

Action

Not available.

Gather Component Alarm Information

Purpose

When you obtain information about component alarms and error messages, you determine when a problem with a component first appeared and the details of the situation.

To gather component alarm information, follow these steps:

1. [Display the Current Router Alarms | 2375](#)
2. [Display Error Messages in the Messages Log File | 2376](#)
3. [Display Error Messages in the Chassis Process Log File | 2376](#)

Display the Current Router Alarms

Purpose

To determine the details of the alarms and when they first appeared in the component.

Action

To display the current router component alarms, use the following CLI command:

```
user@host> show chassis alarms
```


The command output displays the number of alarms currently active, the time when the alarm began, the severity level, and an alarm description. Note the date and time of an alarm so that you can correlate it with error messages in the **messages** system log file.

For examples of sample output, see the *Junos System Basics and Services Command Reference*.

Display Error Messages in the Messages Log File

Purpose

To determine the details of the error messages in the Messages Log File.

Action

To display router component error messages in the **messages** system log file, use the following CLI command:

```
user@host> show log messages
```

The **messages** system log file records the time the failure or event occurred, the severity level, a code, and a message description. Error messages in the **messages** system log file are logged at least 5 minutes before and after the alarm event.

To search for specific information in the log file, use the **| match component-name** command; for example, use **show log messages | match fpc**. If there is a space in the component name, enclose the component name in quotation marks; for example, **| match "power supply"**.

To search for multiple items in the log file, use the **| match** command followed by the multiple items, separated by the **|** (pipe), and enclosed in quotation marks; for example, **| match "fpc | sfm | kernel | tnp"**.

To monitor the **messages** file in real time, use the **monitor start messages** CLI command. This command displays the new entries in the file until you stop monitoring by using the **monitor stop messages** CLI command.

For more information about system log messages, see the *Junos System Log Messages Reference*.

Display Error Messages in the Chassis Process Log File

Purpose

To determine the details of the error messages in the Chassis Process Log File.

Action

To display router component errors in the chassis process (**chassisd**) system log file, use the following CLI command:

```
user@host> show log chassisd
```


The chassis process (**chassisd**) log file tracks the state of each chassis component. For examples of sample output, see the *Junos System Basics and Services Command Reference*.

To search for specific information in the log file, use the **| match component-name** command; for example, **show log messages | match fpc**. If there is a space in the component name, enclose the component name in quotation marks; for example, **| match "power supply"**.

To search for multiple items in the log file, use the **| match** command followed by the multiple items, separated by the **|** (pipe), and enclosed in quotation marks; for example, **| match "fpc | sfm | kernel | tnp"**.

To monitor the **chassisd** file in real time, use the **monitor start chassisd** CLI command. This command displays the new entries in the file until you stop monitoring by using the **monitor stop chassisd** CLI command.

Case Study for a CSPF Failure

Purpose

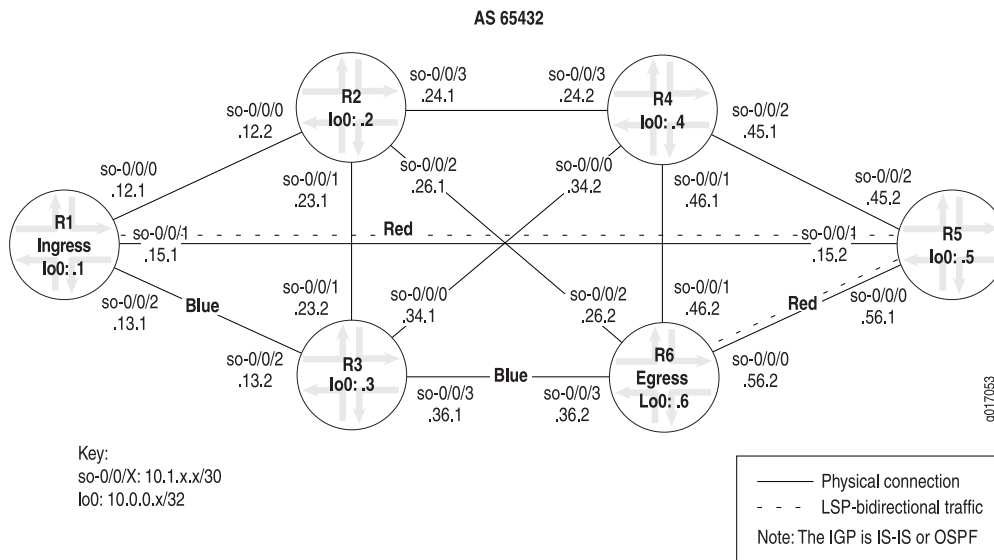
This case study presents a Multiprotocol Label Switching (MPLS) network topology and CSPF failure scenario designed to demonstrate techniques and commands that are particularly useful when addressing CSPF problems in your network. The focus of the study is the incorrect association of user-provided constraints, specifically administrative groups (also known as link coloring).

When calculating a path, the CSPF algorithm factors in user-provided constraints. The ingress router determines the physical path for each LSP by applying a CSPF algorithm to the information in the traffic engineering database. CSPF is a shortest-path-first (SPF) algorithm that has been modified to take into account constraints when calculating the shortest path across the network. Links that do not comply with the restrictions are removed from the tree and cannot be factored into the resulting SPF calculations.

CSPF integrates topology link-state information that is learned from interior gateway protocol (IGP) traffic engineering extensions and is maintained in the traffic engineering database. The information stored in the traffic engineering database includes attributes associated with the state of network resources.

The network topology shown in [Figure 173 on page 2378](#) illustrates a network in which the LSP is constrained by administrative group coloring (also known as link coloring), and CSPF tracing is configured on the ingress router **RI**. In this example, the LSP is forced to transit **R5** in accordance with the restrictions imposed.

Figure 173: CSPF Topology with Administrative Group Coloring



The network shown in [Figure 173 on page 2378](#) is an MPLS router-only network with SONET interfaces.

The MPLS network shown in [Figure 173 on page 2378](#) is configured with administrative group coloring as follows:

- The LSP **R1-to-R6** is established with **R1** as the ingress router and **R6** as the egress router.
- The required path to **R6** transits **R5** on the redlinks. The inclusion of red coloring is not strictly necessary. To force the LSP to transit **R5**, you could color the links on **R3** and **R2** blue and then exclude the blue links.
- Both red and blue colors are used with the **include** and **exclude** statements to ensure that the LSP always transits **R5**. For information on configuring administrative group coloring, see the *MPLS Applications User Guide*.

To check that the network is configured correctly and the LSP is established, follow these steps:

1. [Verify That the LSP Is Established | 2378](#)
2. [Check the Administrative Group Configuration | 2380](#)

Verify That the LSP Is Established

Purpose

Check that the LSP shown in [Figure 173 on page 2378](#) is established and traversing the path from **R1** to **R6** through the red links.

Action

To verify that the LSP is established, enter the following Junos OS command-line interface (CLI) operational mode command:


```
user@host> show mpls lsp extensive
user@host> show mpls lsp
```

Sample Output

```
user@R1> show mpls lsp extensive | no-more
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Metric: 100
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Include: red  Exclude: blue
    Computed ERO
(S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.15.2 S 10.1.56.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.15.2 10.1.56.2
  6 May 11 19:31:42 Selected as active path
  5 May 11 19:31:42 Record Route: 10.1.15.2 10.1.56.2
  4 May 11 19:31:42 Up
  3 May 11 19:31:42 Originate Call
  2 May 11 19:31:42 CSPF: computation result accepted
  1 May 11 19:31:12 CSPF failed: no route toward 10.0.0.6[5 times]
  Created: Wed May 11 19:29:17 2005
Total 1 displayed, Up 1, Down 0
[...Output truncated...]
```

Sample Output 2

```
[edit protocols mpls]
user@R5# run show mpls lsp
```



```
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

```
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Transit LSP: 2 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.0.0.1	10.0.0.6	Up	1	1 FF	100352	3	R6-to-R1
10.0.0.6	10.0.0.1	Up	1	1 FF	100384	3	R1-to-R6

```
Total 2 displayed, Up 2, Down 0
```

Meaning

Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** is successfully established as indicated by the Explicit Route Object (ERO) **10.1.15.2 S 10.1.56.2 S**, the log message **CSPF: computation result accepted**, and **State: Up**. Also, the LSP is routing packets correctly over the red links, avoiding the blue links or the links without any coloring.

Sample Output 2 from transit router **R5** shows that LSP **R1-to-R6** is transiting **R5** as expected.

Check the Administrative Group Configuration

Purpose

Check that the administrative group coloring is correct and the relevant interfaces are associated with each administrative group correctly

Action

To check the administrative group configuration, enter the following Junos OS CLI operational mode commands, or issue the **show** command at the **[edit protocols mpls]** hierarchy level, as shown in the example below:

```
user@host> show configuration protocols mpls
user@host> show mpls interface
user@host> show ted database extensive nodeID
```


Sample Output 1

```
[edit protocols mpls]
user@R1# show
traceoptions {
    file cspf;
    flag cspf;
    flag cspf-node;
    flag cspf-link;
}
admin-groups {
    blue 4;
    red 8;
}
label-switched-path R1-to-R6 {
    to 10.0.0.6;
    metric 100;
    admin-group {
        include red;
        exclude blue;
    }
}
interface so-0/0/0.0;
interface so-0/0/1.0 {
    admin-group red;
}
interface so-0/0/2.0 {
    admin-group blue;
}
interface fxp0.0 {
    disable;
}

[edit protocols mpls]
user@R3# show
admin-groups {
    blue 4;
}
interface fxp0.0 {
    disable;
}
interface so-0/0/0.0;
```



```

interface so-0/0/1.0;
interface so-0/0/2.0 {
interface so-0/0/3.0 {
    admin-group blue;
}

[edit protocols mpls]
user@R5# show
admin-groups {
    red 8;
}
interface fxp0.0 {
    disable;
}
interface so-0/0/0.0 {
    admin-group red;
}
interface so-0/0/1.0;
interface so-0/0/2.0;

[edit protocols mpls]
user@R6# show
admin-groups {
    blue 4;
    red 8;
}
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
interface so-0/0/0.0 {
    admin-group red;
}
interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0 {
    admin-group blue;
}

```


Sample Output 2

```
useruser@R1> show mpls interface
```

```
show mpls interface
```

Interface	State	Administrative groups
so-0/0/0.0	Up	<none>
so-0/0/1.0	Up	red
so-0/0/2.0	Up	blue

```
user@R1> show mpls interface
```

Interface	State	Administrative groups
so-0/0/0.0	Up	<none>
so-0/0/1.0	Up	red
so-0/0/2.0	Up	blue

```
user@R3> show mpls interface
```

Interface	State	Administrative groups
so-0/0/0.0	Up	<none>
so-0/0/1.0	Up	<none>
so-0/0/2.0	Up	<none>
so-0/0/3.0	Up	blue

```
user@R5> show mpls interface
```

Interface	State	Administrative groups
so-0/0/0.0	Up	red
so-0/0/1.0	Up	<none>
so-0/0/2.0	Up	<none>

```
user@R6> show mpls interface
```

Interface	State	Administrative groups
so-0/0/0.0	Up	red
so-0/0/1.0	Up	<none>
so-0/0/2.0	Up	<none>
so-0/0/3.0	Up	blue

Sample Output 3

```

user@R1> show ted database extensive R1
TED database: 6 ISIS nodes 6 INET nodes
NodeID: R1.00(10.0.0.1)
  Type: Rtr, Age: 665 secs, LinkIn: 3, LinkOut: 3
  Protocol: IS-IS(2)
    To: R2.00(10.0.0.2), Local: 10.1.12.1, Remote: 10.1.12.2
      Color: 0 <none>
      Metric: 10
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
        [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
          [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
    To: R5.00(10.0.0.5), Local: 10.1.15.1, Remote: 10.1.15.2
      Color: 0x100 red
      Metric: 10
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
        [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
          [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
    To: R3.00(10.0.0.3), Local: 10.1.13.1, Remote: 10.1.13.2
      Color: 0x10 blue
      Metric: 10
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [priority] bps:
        [0] 155.52Mbps  [1] 155.52Mbps  [2] 155.52Mbps  [3] 155.52Mbps
        [4] 155.52Mbps  [5] 155.52Mbps  [6] 155.52Mbps  [7] 155.52Mbps
      Interface Switching Capability Descriptor(1):

```



```

Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
  [0] 155.52Mbps   [1] 155.52Mbps   [2] 155.52Mbps   [3] 155.52Mbps
  [4] 155.52Mbps   [5] 155.52Mbps   [6] 155.52Mbps   [7] 155.52Mbps

```

Meaning

Sample Output 1 shows that administrative group coloring is correctly configured on all relevant routers. Administrative groups red and blue are configured at the `[edit protocols mpls]` hierarchy level, and relevant interfaces are associated with each administrative group correctly.

R3 is configured with blue coloring and the **include** and **exclude** statements are included in the configuration of **R1** to ensure that LSP **R1-to-R6** always transits **R5**. The inclusion of red coloring is not strictly necessary. To force the LSP to transit **R5**, you could color the links on **R2** and **R3** blue and then exclude the blue links. Red coloring is included in this example to demonstrate the fact that the CSPF algorithm excludes links that do not have a color configured, when the **include** statement is configured at the `[edit protocols mpls]` hierarchy level.

In addition, ingress router **R1** has CSPF tracing configured in preparation for gathering information when the CSPF algorithm fails later in this example.

Sample Output 2 shows that the correct interfaces are associated with the red and blue administration groups on **R1**, **R3**, **R5**, and **R6**.

Sample Output 3 confirms that link coloring is correctly reported in the traffic engineering database for **R1**. Not shown is the traffic engineering database output for the remaining routers, which is similar to the **R1** output, and correct.

Examining a CSPF Failure

When a local CSPF failure indicates that no path meets the constraints configured for the LSP, you must perform CSPF-based tracing and be familiar with the contents of the traffic engineering database to resolve the problem. See [“Examine the Traffic Engineering Database” on page 2390](#) for an analysis of the traffic engineering database.

NOTE: If an LSP does not establish immediately, wait at least a minute or so before taking diagnostic or corrective action. This is because the RSVP retry timer is set to a 30-second default, resulting in a slight delay before the correct state of the LSP is available.

To examine a CSPF failure, follow these steps:

1. [Verify the CSPF Failure | 2386](#)
2. [Examine the CSPF Log File | 2388](#)
3. [Examine the Traffic Engineering Database | 2390](#)
4. [Check the Administrative Group Configuration on R5 | 2393](#)

Verify the CSPF Failure

Purpose

To simulate a configuration error on the network, router **R5** has the administrative group coloring removed from interface **so-0/0/0**. The result is a CSPF failure at **R5** because there is no longer a path between **R1** and **R6** that includes the red color.

Action

To confirm that the LSP is down and verify the configuration on routers **R1** and **R5**, enter the following Junos OS CLI operational mode commands:

```
user@host> clear mpls lsp
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> clear mpls lsp
[edit protocols mpls]
user@R1# run show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 0.0.0.0,   State: Dn
, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Metric: 100
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary                               State: Dn
    Include: red   Exclude: blue
  Will be enqueued for recomputation in 24 second(s).
  9 May 11 20:12:28 CSPF failed: no route toward 10.0.0.6
```



```

8 May 11 20:12:28 Clear Call
7 May 11 20:12:28 Deselected as active
6 May 11 19:31:42 Selected as active path
5 May 11 19:31:42 Record Route: 10.1.15.2 10.1.56.2
4 May 11 19:31:42 Up
3 May 11 19:31:42 Originate Call
2 May 11 19:31:42 CSPF: computation result accepted
1 May 11 19:31:12 CSPF failed: no route toward 10.0.0.6[5 times]
Created: Wed May 11 19:29:17 2005
Total 1 displayed, Up 0, Down 1
[...Output truncated...]

```

Sample Output 2

```

[edit protocols mpls]
user@R5# run show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 1 sessions

```

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.0.0.1	10.0.0.6	Up	1	1 FF	100352	3	R6-to-R1

```

Total 1 displayed, Up 1
, Down 0

```

Meaning

Sample Output 1 from ingress router **R1** shows that the **clear mpls lsp** command was issued to confirm that **R1** cannot reestablish LSP **R1-to-R6**. The sample output from the **show mpls lsp extensive** command shows that LSP **R1-to-R6** is down, **State: Dn** and **ActivePath: (None)**; and that the CSPF has failed, **CSPF failed: no route toward 10.0.0.6**.

Sample Output 2 from transit router **R5** shows that LSP **R1-to-R6** is not included in the output, indicating that the LSP is not transiting **R5**.

Most network problems appear as a local CSPF failure, as shown in the sample output. The CSPF failure indicates that no path meeting the constraints for the LSP can be found in the router's traffic engineering

database. To resolve these problems effectively, use CSPF tracing on the ingress router, and analyze the traffic engineering database to locate the node that should meet the constraints.

Examine the CSPF Log File

Purpose

After you have confirmed that the LSP is down, obtain more information about the possible cause of the failure.

NOTE: To obtain useful information from the CSPF log file, make sure that CSPF tracing is configured on the ingress router.

Action

To examine the CSPF log file, enter the following Junos OS CLI operational mode commands:

```
user@host> monitor start filename
user@host> show log filename
```

NOTE: To stop monitoring CSPF, issue the **monitor stop** command.

Sample Output

```
user@R1> monitor start cspf

[edit protocols mpls]
user@R1# run show log cspf-failed3
May 27 10:22:23 trace_on:   Tracing to "/var/log/cspf"
    started
May 27 10:22:29 CSPF adding path R1-to-R6(primary ) to CSPF queue 1
May 27 10:22:29 CSPF creating CSPF job
May 27 10:22:29
May 27 10:22:29 CSPF for path R1-to-R6(primary ), begin at R1.00 , starting
May 27 10:22:29           path include: 0x00000100
<< administration group red
```



```

May 27 10:22:29          path exclude: 0x00000010
    << administration group blue
May 27 10:22:29          bandwidth: CT0=0bps ; setup priority: 0; random
May 27 10:22:29 CSPF final destination 10.0.0.6
May 27 10:22:29 CSPF starting from R1.00 (10.0.0.1) to 10.0.0.6, hoplimit 254
May 27 10:22:29          constraint include 0x00000100
May 27 10:22:29          constraint exclude 0x00000010
May 27 10:22:29      Node R1.00 (10.0.0.1) metric 0, hops 0, avail 32000 32000 32000
    32000
May 27 10:22:29      Link 10.1.12.1->10.1.12.2(R2.00/10.0.0.2) metric 10 color
    0x00000000 bw 155.52Mbps
May 27 10:22:29      Reverse Link for 10.1.12.1->10.1.12.2 is
    10.1.12.2->10.1.12.1
May 27 10:22:29      link fails include 0x00000100
May 27 10:22:29      Link 10.1.15.1->10.1.15.2(R5.00/10.0.0.5) metric 10 color
    0x00000100 bw 155.52Mbps
May 27 10:22:29      Reverse Link for 10.1.15.1->10.1.15.2 is
    10.1.15.2->10.1.15.1
May 27 10:22:29      link's interface switch capability descriptor #1
May 27 10:22:29      encoding: Packet, switching: Packet
May 27 10:22:29      link passes constraints
May 27 10:22:29      Link 10.1.13.1->10.1.13.2(R3.00/10.0.0.3) metric 10 color
    0x00000010 bw 155.52Mbps
May 27 10:22:29      Reverse Link for 10.1.13.1->10.1.13.2 is
    10.1.13.2->10.1.13.1
May 27 10:22:29      link fails include 0x00000100
May 27 10:22:29      Node R5.00 (10.0.0.5) metric 10, hops 1, avail 32000 32000 32000
    32000
May 27 10:22:29      Link 10.1.15.2->10.1.15.1(R1.00/10.0.0.1) metric 10 color
    0x00000100 bw 155.52Mbps
May 27 10:22:29      skipped: end point already visited
May 27 10:22:29      Link 10.1.45.2->10.1.45.1(R4.00/10.0.0.4) metric 10 color
    0x00000000 bw 155.52Mbps
May 27 10:22:29      Reverse Link for 10.1.45.2->10.1.45.1 is
    10.1.45.1->10.1.45.2
May 27 10:22:29      link fails include 0x00000100
May 27 10:22:29      Link 10.1.56.1->10.1.56.2(R6.00/10.0.0.6) metric 10 color
    0x00000000 bw 155.52Mbps
May 27 10:22:29      Reverse Link for 10.1.56.1->10.1.56.2 is 10.1.56.2->10.1.56.1
May 27 10:22:29      link fails include 0x00000100
May 27 10:22:29 CSPF completed in 0s
May 27 10:22:29 CSPF couldn't find a route to 10.0.0.6
May 27 10:22:29 CSPF for R1-to-R6 done!
monitor stop

```


Meaning

The sample output shows that the **monitor start cspf** command was issued to start displaying entries in the **cspf** log file in real time. The **cspf** log file is generated by the routing protocol process after the file is configured with the **traceoptions** statement at the **[edit protocols mpls]** hierarchy level. In this example, the **cspf** log file is configured with the **cspf**, **cspf-node**, and **cspf-link** flags to provide the most granular information about the steps taken by the CSPF algorithm.

The only link that passes the color constraint is between **R1** and **R5**, **10.1.15.0/32**. The CSPF algorithm is a locally run algorithm, which makes its calculations on a given router. When the CSPF algorithm runs on **R5**, it prunes **10.1.15.2** and selects **10.1.56.1** to send the message to **R6**. The link between **R5** and **R6** **10.1.56.0/32** does not pass the color constraints, indicating a problem with **R5**. At this stage, it is useful to examine the traffic engineering database to determine which link on **R5** should be associated with the red color.

Examine the Traffic Engineering Database

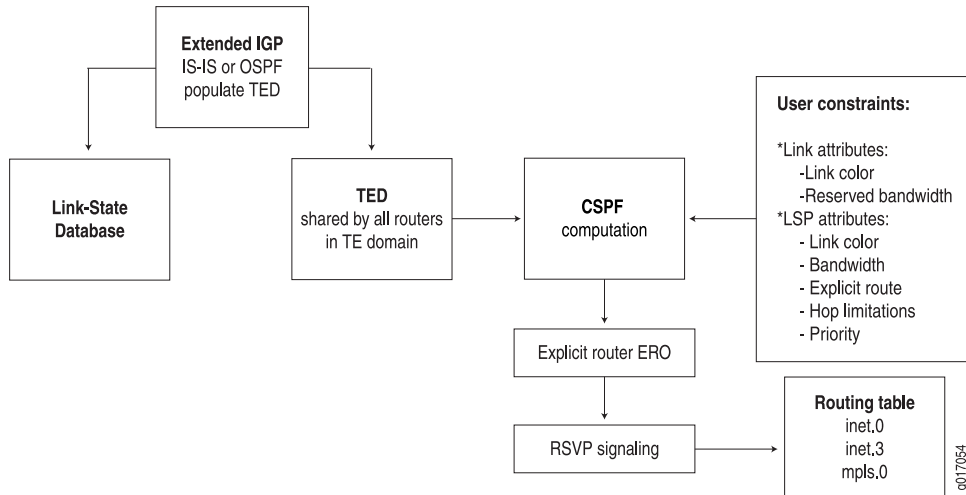
Purpose

Examining the traffic engineering database is another way to locate the node that should meet the constraints but does not. Once identified, you can concentrate your troubleshooting efforts on why that node is not being represented accurately in the database.

The contents of the traffic engineering database are consistent among all routers within a given traffic engineering domain. Therefore, you can issue the **show ted database** command from any router in the same traffic engineering domain to obtain more granular information about the CSPF failure.

CSPF integrates topology link-state information that is learned from the IGP traffic engineering extensions and maintained in the traffic engineering database. The information stored in the traffic engineering database includes attributes associated with the state of the network resources (such as total link bandwidth, reserved link bandwidth, available link bandwidth, and link color). When calculating a path, the CSPF algorithm factors in user-provided information such as bandwidth requirements, maximum allowed hop count, and administrative groups, all of which are obtained from user configuration. (See [Figure 174 on page 2391](#)).

Figure 174: User-Provided Constraints



Action

To examine the traffic engineering database, enter the following Junos OS CLI operational mode commands:

```

user@host> show ted database extensive
user@host> show ted database extensive NodeID | match "(NodeID|To:|Color)"

```

Sample Output 1

```

[edit protocols mpls]
user@R1# run show ted database extensive
TED database: 6 ISIS nodes 6 INET nodes
[...Output truncated...]
NodeID: R5.00(10.0.0.5)
  Type: Rtr , Age: 103 secs, LinkIn: 3, LinkOut: 3
  Protocol: IS-IS(2)
  To: R1.00(10.0.0.1), Local: 10.1.15.2, Remote: 10.1.15.1
  Color: 0x100 red
  Metric: 10
  Static BW: 155.52Mbps
  Reservable BW: 155.52Mbps
  Available BW [priority] bps:
    [0] 155.52Mbps [1] 155.52Mbps [2] 155.52Mbps [3] 155.52Mbps
    [4] 155.52Mbps [5] 155.52Mbps [6] 155.52Mbps [7] 155.52Mbps
  Interface Switching Capability Descriptor(1):

```



```

Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
    [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
To: R4.00(10.0.0.4) , Local: 10.1.45.2, Remote: 10.1.45.1
Color: 0 <none>
Metric: 10
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
    [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
    [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
    [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
To: R6.00(10.0.0.6), Local: 10.1.56.1, Remote: 10.1.56.2
Color: 0 <none>
Metric: 10
Static BW: 155.52Mbps
Reservable BW: 155.52Mbps
Available BW [priority] bps:
    [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
    [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
Interface Switching Capability Descriptor(1):
Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 155.52Mbps    [1] 155.52Mbps    [2] 155.52Mbps    [3] 155.52Mbps
    [4] 155.52Mbps    [5] 155.52Mbps    [6] 155.52Mbps    [7] 155.52Mbps
[...Output truncated...]

```

Sample Output 2

```

[edit protocols]
user@R1# run show ted database extensive R5.00 | match "(NodeID|To:|Color)"
NodeID: R5.00(10.0.0.5)

```



```

To: R1.00(10.0.0.1), Local: 10.1.15.2, Remote: 10.1.15.1
  Color: 0x100 red
To: R4.00(10.0.0.4), Local: 10.1.45.2, Remote: 10.1.45.1
  Color: 0 <none>
To: R6.00(10.0.0.6), Local: 10.1.56.1, Remote: 10.1.56.2
  Color: 0 <none>
To: R1.00(10.0.0.1), Local: 10.1.15.2, Remote: 10.1.15.1
  Color: 0x100 red
To: R4.00(10.0.0.4), Local: 10.1.45.2, Remote: 10.1.45.1
  Color: 0 <none>
To: R6.00(10.0.0.6), Local: 10.1.56.1, Remote: 10.1.56.2
  Color: 0 <none>

```

Meaning

Sample Output 1 from ingress router **R1** shows a wealth of information on each node in the network, although only a portion is included in this example. The output shows the total number of IS-IS and INET nodes in the traffic engineering domain. The portion of the traffic engineering database shown represents a node (**R5**), and the **Type** field indicates **Rtr** (router). The **Type** field could also indicate Net (network) if the node were a pseudo node. The node (**R5**) has three input and output links that are running IS-IS Level 2, **Protocol: IS-IS(2)**. The links lead to nodes **R1**, **R4**, and **R6**. The local address and remote address for each link is specified. The information on each node includes administrative groups (**Color:**), metrics, static bandwidth, reservable bandwidth, and available bandwidth priority level. The information contained in the traffic engineering database should be the same across all routers in the same traffic engineering domain. For a detailed description of the fields in the output of the **show ted database extensive** command, see the *Junos Routing Protocols and Policies Command Reference*.

Sample Output 2 shows filtered output that allows you to focus on exactly what is missing or incorrect.

Both outputs confirm that the link between **R1** and **R5**, **10.1.15.0/32**, is associated with the red color, while the link between **R5** and **R6**, **10.1.56.0/32**, is not associated with a color. In the network shown in [Figure 173 on page 2378](#), for the LSP to establish correctly, link **10.1.56.1** must be associated with the red color.

Check the Administrative Group Configuration on R5

Purpose

Focus on R5 to determine which interfaces are associated with the red color, and make any necessary corrections.

Action

To check the administrative group configuration on R5 and make any necessary corrections, enter the following Junos OS CLI commands:


```

user@R5> edit
[edit protocols mpls]
user@R5# show
user@R5# delete interface so-0/0/1 admin-group
user@R5# set interface so-0/0/0 admin-group red
user@R5# show
user@R5# commit

```

Sample Output 1

```

user@R5> edit
Entering configuration mode

[edit protocols mpls]
user@R5# show
admin-groups {
    red 8;
}
interface fxp0.0 {
    disable;
}
interface so-0/0/0.0;
interface so-0/0/1.0 { <<<incorrect interface configured with admin-group
    admin-group red;
}
interface so-0/0/2.0;

```

Sample Output 2

```

[edit protocols mpls]
user@R5# delete interface so-0/0/1 admin-group

[edit protocols mpls]
user@R5# set interface so-0/0/0 admin-group red

[edit protocols mpls]

```



```

user@R5# show
admin-groups {
    red 8;
    blue 4;
}
interface fxp0.0 {
    disable;
}
interface so-0/0/0.0 { <<<correct interface configured with admin-group
    admin-group red;
}
interface so-0/0/1.0;
interface so-0/0/2.0;

[edit protocols mpls]
user@R5# commit
commit complete

```

Sample Output 3

```

user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Up    1      * R1-to-R6
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPname
10.0.0.1    10.0.0.6    Up    0  1 FF      3      - R6-to-R1
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning

Sample Output 1 from transit router **R5** shows that at the `[edit protocols mpls]` hierarchy level, interface **so-0/0/1** is incorrectly configured with the **admin-group red** statement. The **so-0/0/0** interface should be configured with the **admin-group red** statement.

Sample Output 2 shows the steps taken to correct the configuration. The administration group has been deleted from **so-0/0/1** and **so-0/0/0** is now associated with the red color.

Sample Output 3 shows that LSP **R1-to-R6** is established.

Verify the CSPF Failure

Purpose

To simulate a configuration error on the network, router **R5** has the administrative group coloring removed from interface **so-0/0/0**. The result is a CSPF failure at **R5** because there is no longer a path between **R1** and **R6** that includes the red color.

Action

To confirm that the LSP is down and verify the configuration on routers **R1** and **R5**, enter the following Junos OS CLI operational mode commands:

```
user@host> clear mpls lsp
user@host> show mpls lsp extensive
```

Sample Output 1

```
user@R1> clear mpls lsp
[edit protocols mpls]
user@R1# run show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 0.0.0.0, State: Dn
, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
LoadBalance: Random
Metric: 100
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary State: Dn
  Include: red Exclude: blue
Will be enqueued for recomputation in 24 second(s).
  9 May 11 20:12:28 CSPF failed: no route toward 10.0.0.6
  8 May 11 20:12:28 Clear Call
  7 May 11 20:12:28 Deselected as active
```



```

6 May 11 19:31:42 Selected as active path
5 May 11 19:31:42 Record Route: 10.1.15.2 10.1.56.2
4 May 11 19:31:42 Up
3 May 11 19:31:42 Originate Call
2 May 11 19:31:42 CSPF: computation result accepted
1 May 11 19:31:12 CSPF failed: no route toward 10.0.0.6[5 times]
Created: Wed May 11 19:29:17 2005
Total 1 displayed, Up 0, Down 1
[...Output truncated...]

```

Sample Output 2

```

[edit protocols mpls]
user@R5# run show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 1 sessions

```

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.0.0.1	10.0.0.6	Up	1	1 FF	100352	3	R6-to-R1

```

Total 1 displayed, Up 1
, Down 0

```

Meaning

Sample Output 1 from ingress router **R1** shows that the **clear mpls lsp** command was issued to confirm that **R1** cannot reestablish LSP **R1-to-R6**. The sample output from the **show mpls lsp extensive** command shows that LSP **R1-to-R6** is down, **State: Dn** and **ActivePath: (None)**; and that the CSPF has failed, **CSPF failed: no route toward 10.0.0.6**.

Sample Output 2 from transit router **R5** shows that LSP **R1-to-R6** is not included in the output, indicating that the LSP is not transiting **R5**.

Most network problems appear as a local CSPF failure, as shown in the sample output. The CSPF failure indicates that no path meeting the constraints for the LSP can be found in the router's traffic engineering database. To resolve these problems effectively, use CSPF tracing on the ingress router, and analyze the traffic engineering database to locate the node that should meet the constraints.

Examining the Hello Message

Purpose

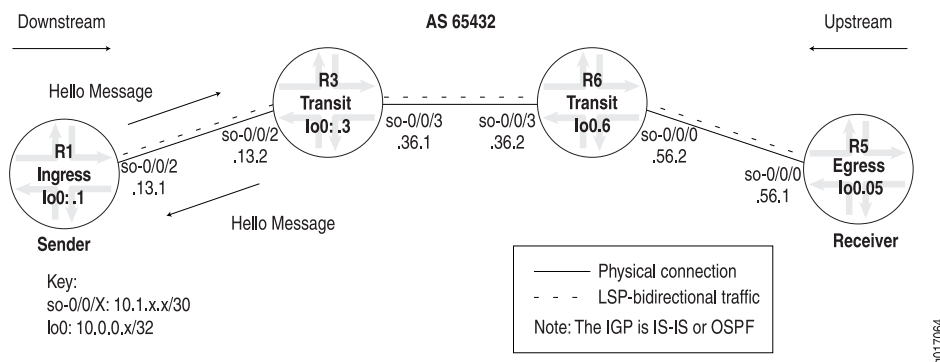
RSVP monitors the status of the interior gateway protocol (IGP) (Intermediate System-to-Intermediate System [ISIS] or Open Shortest Path First [OSPF]) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because Hello messages are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

RSVP Hello messages are exchanged between neighbors. The destination address is the neighbor node. RSVP Hello messages are used to determine loss of interface more quickly than determined by the RSVP state timeout.

NOTE: RSVP Hello messages are required to establish the protocol or to maintain adjacency information. RSVP Hello messages do not establish state.

Figure 175 on page 2398 shows two RSVP Hello messages exchanged between routers R1 and R3.

Figure 175: RSVP Hello Message



To ensure that Hello messages are displayed in the output, include the **packets** flag at the **[edit protocols rsvp traceoptions]** hierarchy level.

Action

To examine the Hello message, enter the following Junos OS CLI command:

```
user@R1> monitor start filename
```


Sample Output 1

```
[edit protocols rsvp]
user@R1# show
traceoptions {
    file rsvp-log;
    flag packets detail;
}
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}
```

Sample Output 2

```
user@R1> clear log rsvp-log

user@R1> monitor start rsvp-log

user@R1>
*** rsvp-log ***
[...Output truncated...]
Jun 29 15:48:59  RSVP send Hello New 10.1.13.1->10.1.13.2 Len=32 so-0/0/2.0
Jun 29 15:48:59  HelloReq Len 12
Jun 29 15:48:59  RestartCap Len 12 restart time 0, recovery time 0
Jun 29 15:48:59  RSVP rcv Hello New 10.1.13.2->10.1.13.1 Len=32 so-0/0/2.0
Jun 29 15:48:59  HelloRply Len 12
Jun 29 15:48:59  RestartCap Len 12 restart time 0, recovery time 0
monitor stop
```

Meaning

Sample Output 1 shows the configuration of RSVP tracing on ingress router **R1**. The **packets** flag is included at the **[edit protocols rsvp traceoptions]** hierarchy level to provide information about RSVP traffic. The **detail** option is included to show granular details about the configured flag.

Sample Output 2 shows **clear** commands, the output for the **rsvp-log** file, and that monitoring was started and then stopped. The **rsvp-log** output shows two RSVP Hello messages exchanged between **R1** and **R3**.

The first Hello message in the **rsvp-log** output is a request sent from **R1 (10.1.13.1)** to **R3 (10.1.13.2)**. The outgoing interface is **so-0/0/2.0** on **R1**. The second Hello message was a reply sent from **R3** to **R1**, also through the outgoing interface **so-0/0/2.0** on **R3**.

The next two lines of output indicate object values for the two Hello messages, and are indented in the output. To facilitate this discussion, each line of output for each object is displayed before the corresponding explanation.

- **HelloReq Len 12**

The Hello request (**HelloReq**) object indicates that this is a Hello request. RFC 3209 defines the RSVP Hello message. An RSVP Hello message can either be a request or a reply. Every request should generate a reply.

- **RestartCap Len 12 restart time 0, recovery time 0**

The restart object (**RestartCap**) indicates the graceful restart capability of the sender node. The restart time of 0 milliseconds is the length of time that this node takes to restart its RSVP traffic engineering functionality. At the end of this time, the node can send and receive RSVP messages again. The recovery time of 0 milliseconds indicates the length of time the LSR retains MPLS forwarding information. A recovery time of 0 in this case indicates that no forwarding state was preserved across a restart. Because both values are set to 0, graceful restart was not enabled for this RSVP session.

- **HelloRply Len 12**

The Hello reply (**HelloRply**) object indicates that this is an RSVP Hello message sent from **R3** to **R1** out of interface **so-0/0/2.0**.

In standard RSVP, node failure detection occurs as a consequence of the RSVP soft-state timeout model. However, detection typically requires several minutes to time out the soft state. Hello packets allow the detection of the neighboring node state changes more quickly.

In Junos OS, RSVP Hello messages are optional and are backward-compatible with RSVP implementations that do not support Hello messages. For neighboring routers that do not support Hello messages or on which RSVP Hello is disabled, RSVP uses the soft-state timeout for loss detection and cannot benefit from fast IGP Hello detection.

Configuring a short time for the IS-IS or OSPF Hello timers allows these protocols to detect node failures more quickly. RSVP also benefits from early detection by the IGP protocols. It is not necessary to explicitly configure a short RSVP Hello timer. If you do configure the RSVP Hello timer, you can configure a longer value and can still expect the failure of a neighboring router to be quickly detected by IGP.

Between Hello-capable neighbors, Hello messages are sent unicast toward each other. A loss of (2 x keep-multiplier + 1) consecutive Hello messages causes the neighbor's state to go down, and all RSVP sessions to and from that neighbor are declared to be down.

By default, RSVP sends Hello messages every 9 seconds.

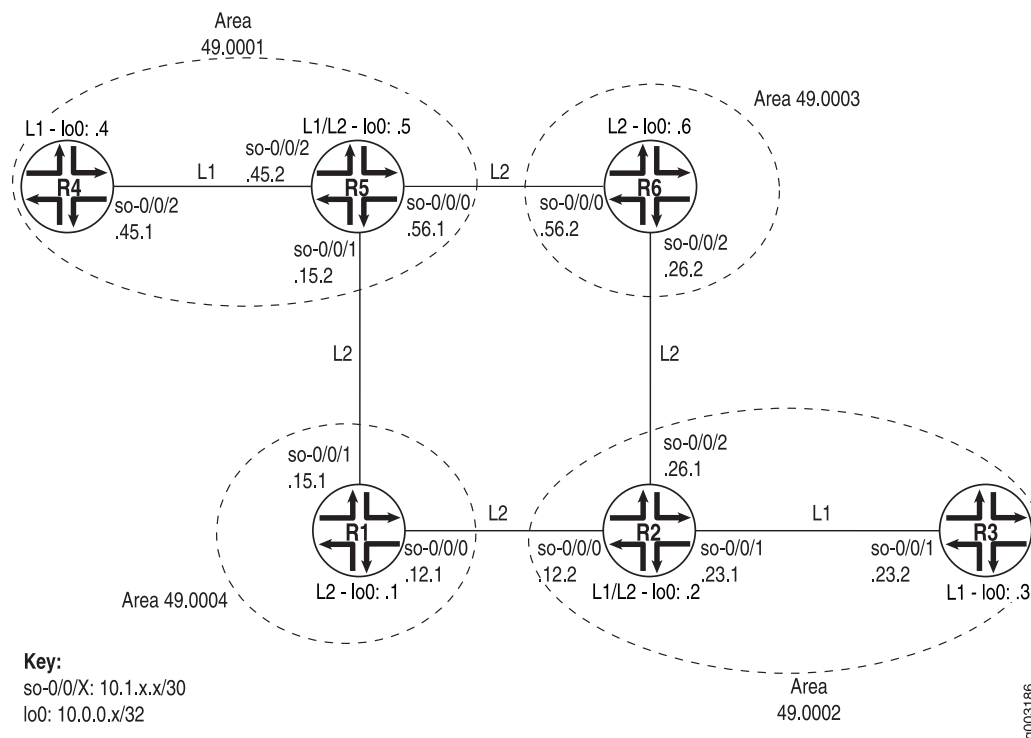
Displaying the Status of IS-IS Adjacencies

Purpose

Assuming that all the routers are correctly configured for IS-IS, you can verify which neighbors are adjacent and able to exchange IS-IS data. In addition, you can examine the set of routes installed in the forwarding table to verify that the routing protocol process (rpd) has relayed the correct information into the forwarding table.

Figure 176 on page 2401 illustrates the example IS-IS topology used for the procedures in this topic.

Figure 176: IS-IS Network Topology



The network consists of Level 1 and Level 2 adjacencies. Level 1 adjacencies are within areas 49.0001 and 49.0002. Level 2 adjacencies occur between all directly connected Level 2 routers regardless of which area they are in. For example, R5 is in area 49.0001, R6 is in area 49.0003, R1 is in area 49.0004, and R2 is in area 49.0002. The network in Figure 176 on page 2401 should have the following adjacencies:

- Level 2 adjacencies between all directly connected Level 2 routers (R1, R2, R5, and R6).
- Level 1 adjacencies between routers in area 49.0001 (R4 and R5) and between routers in area 49.0002 (R2 and R3).

To verify that routers are adjacent and able to exchange IS-IS data, follow these steps:

1. [Verifying Adjacent Routers | 2402](#)
2. [Examine the Forwarding Table | 2403](#)

Verifying Adjacent Routers

Purpose

Verify that routers are adjacent and able to exchange IS-IS data.

Action

To verify that routers are adjacent and able to exchange IS-IS data, enter the following CLI operational mode command:

```
user@host> show isis adjacency
```

The following sample output shows the adjacencies that formed for all routers shown in [“Displaying the Status of IS-IS Adjacencies” on page 2401](#).

Sample Output

```
user@R1> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R2          2 Up         19
so-0/0/1.0         R5          2 Up         18

user@R2> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R1          2 Up         19
so-0/0/1.0         R3          1 Up         26
so-0/0/2.0         R6          2 Up         21

user@R3> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/1.0         R2          1 Up         24

user@R4> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/2.0         R5          1 Up         23
```



```

user@R5> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R6          2 Up         22
so-0/0/1.0         R1          2 Up         20
so-0/0/2.0         R4          1 Up         20

user@R6> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
so-0/0/0.0         R5          2 Up         21
so-0/0/2.0         R2          2 Up         20

```

Meaning

The sample output shows the adjacencies that formed in the network illustrated in [“Displaying the Status of IS-IS Adjacencies” on page 2401](#). The Level 1/Level 2 routers (R2 and R5) formed Level 1 adjacencies with Level 1 routers (R3 and R4), and Level 2 adjacencies with the Level 2 routers (R1 and R6). To view the status of the adjacency, examine the State column. In this example, all adjacencies in the network are up.

If the state is not **Up** for a particular neighbor, you must first examine the IS-IS configuration for the particular interface. Make sure that the NET address is correct and that the loopback interface (lo0) is configured. Use the **show isis interface** or **show isis interface detail** command to display the IS-IS parameters for all interfaces configured with IS-IS. With these two commands, you can see which interfaces are configured for IS-IS, whether they are configured for Level 1 or Level 2, the IS-IS metric, and other IS-IS information.

SEE ALSO

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

Examine the Forwarding Table

Purpose

You can display the set of routes installed in the forwarding table to verify that the routing protocol process (rpd) has relayed the correct information into the forwarding table. This is especially important when there are network problems, such as connectivity. In this procedure, you verify that the routes displayed in Step 2 appear in the forwarding table for Router R5.

Action

To examine the forwarding table for a router, enter the following CLI command:


```
user@host> show route forwarding-table destination destination-prefix
```

Sample Output

```
user@R5> show route forwarding-table destination 10.0.0.3
```

```
Routing table: inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
10.0.0.3/32          user    0 10.1.15.0          ucst  285    7  so-0/0/1.0

user@R5> show route forwarding-table destination 10.0.0.3
Routing table: inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
10.0.0.3/32          user    0 10.1.56.0          ucst  281    9  so-0/0/0.0
```

Meaning

The sample output shows the selected next hop between Routers R5 and R3 sent from the inet routing table and installed into the forwarding table. The first instance shows the route through Router R1, and the second instance shows the route through Router R6. In both instances, the preferred route displayed in Step 2 is installed in the forwarding table.

In general, the sample output includes the destination address and destination type, the next-hop address and next-hop type, the number of references to the next hop, an index number into an internal next-hop database, and the interface used to reach the next hop.

SEE ALSO

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

RELATED DOCUMENTATION

Verifying the IS-IS Protocol

Check OSPF on a Stub Router

Purpose

To verify the OSPF configuration on a stub router.

Action

To verify the OSPF configuration on a stub router in your network, enter the following commands:

```
user@host> show configuration
user@host> show ospf interface
```

The following sample output is for an OSPF configuration on **R5**, a stub router:

Sample Output

```
user@R5> show configuration
[...Output truncated...]
interfaces {
    so-0/0/2 {
        unit 0 {
            family inet {
                address 10.1.45.2/30;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.5/32;
            }
        }
    }
}
routing-options {
    router-id 10.0.0.5;
}
protocols {
    ospf {
        area 0.0.0.2 {
            stub;
        }
    }
}
```



```

        interface so-0/0/2.0;
        interface lo0.0 {
            passive;
        }
    }
}

```

```

user@R5> show ospf interface

```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DRother	0.0.0.2	0.0.0.0	0.0.0.0	0
so-0/0/2.0	PtToPt	0.0.0.2	0.0.0.0	0.0.0.0	1

Meaning

The sample output shows a basic OSPF configuration at the [edit protocols ospf] and [edit interfaces] hierarchy levels on **R5**, a stub router.

R5 has two interfaces included at the [edit protocols ospf] hierarchy level, and those interfaces have the **family inet** statement included at the [edit interfaces] hierarchy level. Both interfaces, **so-0/0/2.0** and the loopback interface (**lo0**), are in the stub area (**0.0.0.2**).

R5 has the router ID configured manually to avoid possible problems when the OSPF router ID (RID) changes; for example, when multiple loopback addresses are configured. The RID uniquely identifies the router within the OSPF network. It is transmitted within the LSAs used to populate the link-state database and calculate the shortest-path tree. In a link-state network, it is important that two routers do not share the same RID value, otherwise IP routing problems may occur.

A stub area does not allow AS external advertisements to flood within that area. **R5** relies on a default route (**0.0.0.0/0**) to reach destinations outside the AS. The default route can be statically configured on **R5** or advertised by an ABR (**R4**). In this network, the default LSA is advertised by **R4**.

A stub area is useful if you want to reduce the size of the topological database and therefore the amount of memory required from the routers in the stub area. However, some restrictions apply to a stub area. You cannot create a virtual link through a stub area, and a stub area cannot contain an ASBR.

Checklist for Verifying the BGP Protocol and Peers

Purpose

[Table 69 on page 2407](#) provides links and commands for verifying whether the Border Gateway Protocol (BGP) is configured correctly on a Juniper Networks router in your network, the internal Border Gateway Protocol (IBGP) and exterior Border Gateway Protocol (EBGP) sessions are properly established, the

external routes are advertised and received correctly, and the BGP path selection process is working properly.

Action

Table 69: Checklist for Verifying the BGP Protocol and Peers

Tasks	Command or Action
“Verify BGP Peers” on page 2407	
1. <i>Verify BGP on an Internal Router</i>	<code>show configuration</code>
2. <i>Verify BGP on a Border Router</i>	<code>show configuration</code>
2. <i>Verify Advertised BGP Routes</i>	<code>show route advertising-protocol bgp neighbor-address</code>
3. <i>Verify That a Particular BGP Route Is Received on Your Router</i>	<code>show route receive-protocol bgp neighbor-address</code>
“Examine BGP Routes and Route Selection” on page 2418	
1. Examine the Local Preference Selection on page 2420	<code>show route destination-prefix < detail ></code>
2. Examine the Multiple Exit Discriminator Route Selection on page 2422	<code>show route destination-prefix < detail ></code>
3. Examine the EBGp over IBGP Selection on page 2423	<code>show route destination-prefix < detail ></code>
4. Examine the IGP Cost Selection on page 2425	<code>show route destination-prefix < detail ></code>
“Examine Routes in the Forwarding Table” on page 2432	<code>show route forwarding-table</code>

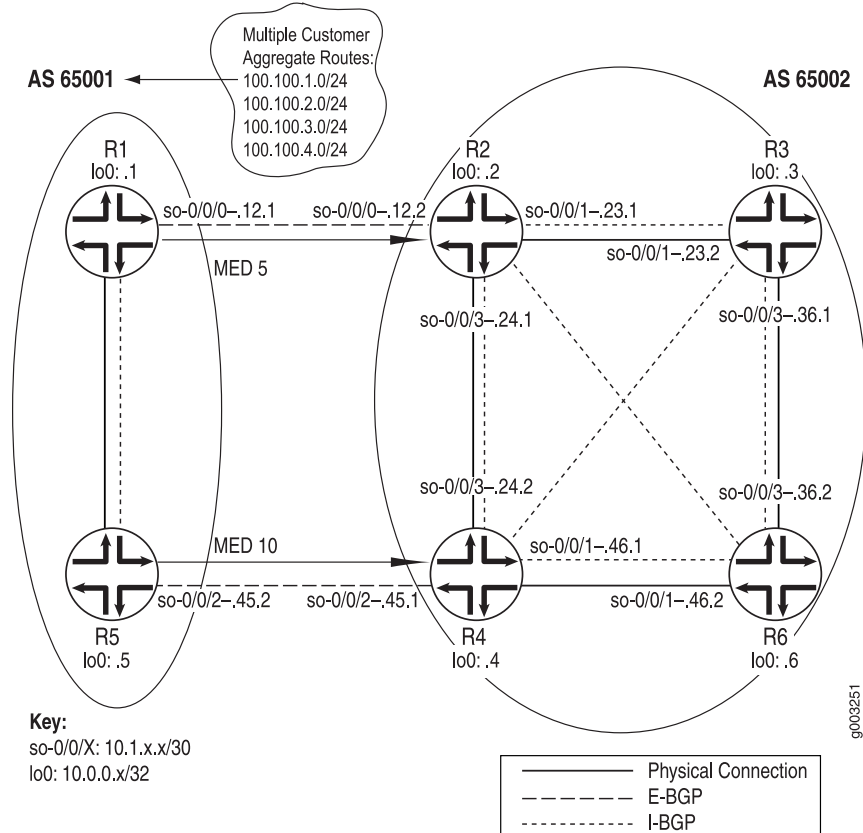
Verify BGP Peers

Purpose

Assuming that all the routers are correctly configured for BGP, you can verify if IBGP and EBGp sessions are properly established, external routes are advertised and received correctly, and the BGP path selection process is working properly.

[Figure 177 on page 2408](#) illustrates an example BGP network topology used in this topic.

Figure 177: BGP Network Topology



The network consists of two directly connected ASs consisting of external and internal peers. The external peers are directly connected through a shared interface and are running EBGp. The internal peers are connected through their loopback (lo0) interfaces through IBGP. AS 65001 is running OSPF and AS 65002 is running IS-IS as its underlying IGP. IBGP routers do not have to be directly connected, the underlying IGP allows neighbors to reach one another.

The two routers in AS 65001 each contain one EBGp link to AS 65002 (R2 and R4) over which they announce aggregated prefixes: **100.100.1.0**, **100.100.2.0**, **100.100.3.0**, and **100.100.4.0**. Also, R1 and R5 are injecting multiple exit discriminator (MED) values of 5 and 10, respectively, for some routes.

The internal routers in both ASs are using a full mesh IBGP topology. A full mesh is required because the networks are not using confederations or route reflectors, so any routes learned through IBGP are not distributed to other internal neighbors. For example, when R3 learns a route from R2, R3 does not distribute that route to R6 because the route is learned through IBGP, so R6 must have a direct BGP connection to R2 to learn the route.

In a full mesh topology, only the border router receiving external BGP information distributes that information to other routers within its AS. The receiving router does not redistribute that information to other IBGP routers in its own AS.

From the point of view of AS 65002, the following sessions should be up:

- The four routers in AS 65002 should have IBGP sessions established with each other.
- R2 should have an EBGP session established with R1.
- R4 should have an EBGP session established with R5.

To verify BGP peers, follow these steps:

1. [Verify BGP on an Internal Router | 2409](#)
2. [Verify BGP on a Border Router | 2412](#)
3. [Verify Advertised BGP Routes | 2417](#)
4. [Verify That a Particular BGP Route Is Received on Your Router | 2417](#)

Verify BGP on an Internal Router

Purpose

To verify the BGP configuration of an internal router.

Action

To verify the BGP configuration of an internal router, enter the following Junos OS command-line interface (CLI) command:

```
user@host> show configuration
```

The following sample output is for a BGP configuration on R3:

Sample Output

```
user@R3> show configuration
[...Output truncated...]
interfaces {
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.23.2/30;
      }
      family iso;
    }
  }
  so-0/0/3 {
    unit 0 {
```



```

        family inet {
            address 10.1.36.1/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.3/32;
        }
        family iso {
            address 49.0002.1000.0000.0003.00;
        }
    }
}
}
routing-options {
    [...Output truncated...]
    router-id 10.0.0.3;
    autonomous-system 65002;
}
protocols {
    bgp {
        group internal {
            type internal;
            local-address 10.0.0.3;
            neighbor 10.0.0.2;
            neighbor 10.0.0.4;
            neighbor 10.0.0.6;
        }
    }
    isis {
        level 1 disable;
        interface all {
            level 2 metric 10;
        }
        interface lo0.0;
    }
}

```

```
user@R6> show configuration |
```



```

[Output truncated...]
interfaces {
    so-0/0/1 {
        unit 0 {
            family inet {
                address 10.1.46.2/30;
            }
            family iso;
        }
    }
    so-0/0/3 {
        unit 0 {
            family inet {
                address 10.1.36.2/30;
            }
            family iso;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 10.0.0.6/32;
            }
            family iso {
                address 49.0003.1000.0000.0006.00;
            }
        }
    }
}
routing-options {
    [Output truncated...]
    router-id 10.0.0.6;
    autonomous-system 65002;
}
protocols {
    bgp {
        group internal {
            type internal;
            local-address 10.0.0.6;
            neighbor 10.0.0.2;
            neighbor 10.0.0.3;
            neighbor 10.0.0.4;

```



```

    }
  }
  isis {
    level 1 disable;
    interface all {
      level 2 metric 10;
    }
    interface lo0.0;
  }
}

```

Meaning

The sample output shows a basic BGP configuration on routers **R3** and **R6**. The local AS (65002) and one group (**internal**) are configured on both routers. **R3** has three internal peers—**10.0.0.2**, **10.0.0.4**, and **10.0.0.6**—included at the `[protocols bgp group group]` hierarchy level. **R6** also has three internal peers: **10.0.0.2**, **10.0.0.3**, and **10.0.0.4**. The underlying IGP protocol is Intermediate System-to-Intermediate System (IS-IS), and relevant interfaces are configured to run IS-IS.

Note that in this configuration the router ID is manually configured to avoid any duplicate router ID problems.

Verify BGP on a Border Router

Purpose

To verify the BGP configuration of a border router.

Action

To verify the BGP configuration of a border router, enter the following Junos OS CLI operational mode command:

```
user@host> show configuration
```

Sample Output

The following sample output is for a BGP configuration on two border routers, R2 and R4 from AS 65002:

```

user@R2> show configuration
[...Output truncated...]
interfaces {

```



```

so-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.12.2/30;
        }
        family iso;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.23.1/30;
        }
        family iso;
    }
}
so-0/0/3 {
    unit 0 {
        family inet {
            address 10.1.24.1/30;
        }
        family iso;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.0.0.2/32;
        }
        family iso {
            address 49.0002.1000.0000.0002.00;
        }
    }
}
}
routing-options {
    [...Output truncated...]
    router-id 10.0.0.2;
    autonomous-system 65002;
}
protocols {
    bgp {
        group internal {

```



```

        type internal;
        export next-hop-self;
        neighbor 10.0.0.3;
        neighbor 10.0.0.4;
        neighbor 10.0.0.6;
    }
    group toR1 {
        type external;
        import import-toR1;
        peer-as 65001;
        neighbor 10.1.12.1;
    }
}
isis {
    level 1 disable;
    interface all {
        level 2 metric 10;
    }
    interface lo0.0;
}
}
policy-options {
    policy-statement next-hop-self {
        term change-next-hop {
            from neighbor 10.1.12.1;
            then {
                next-hop self;
            }
        }
    }
    policy-statement import-toR1 {
        term 1 {
            from {
                route-filter 100.100.1.0/24 exact;
            }
            then {
                local-preference 200;
            }
        }
    }
}

```

user@R4> **show configuration**

[...Output truncated...]


```

interfaces {
  so-0/0/1 {
    unit 0 {
      family inet {
        address 10.1.46.1/30;
      }
      family iso;
    }
  }
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.1.45.1/30;
      }
      family iso;
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.24.2/30;
      }
      family iso;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.4/32;
      }
      family iso {
        address 49.0001.1000.0000.0004.00;
      }
    }
  }
}

routing-options {
  [...Output truncated...]
  router-id 10.0.0.4;
  autonomous-system 65002;
}

protocols {
  bgp {

```



```

    group internal {
        type internal;
        local-address 10.0.0.4;
        export next-hop-self;
        neighbor 10.0.0.2;
        neighbor 10.0.0.3;
        neighbor 10.0.0.6;
    }
    group toR5 {
        type external;
        peer-as 65001;
        neighbor 10.1.45.2;
    }
}
isis {
    level 1 disable;
    interface all {
        level 2 metric 10;
    }
    interface lo0.0;
}
}
policy-options {
    policy-statement next-hop-self {
        term change-next-hop {
            from neighbor 10.1.45.2;
            then {
                next-hop self;
            }
        }
    }
}

```

Meaning

The sample output shows a basic BGP configuration on border routers **R2** and **R4**. Both routers have the AS (65002) included at the **[routing-options]** hierarchy level. Each router has two groups included at the **[protocols bgp group group]** hierarchy level. External peers are included in the external group, either **toR1** or **toR5**, depending on the router. Internal peers are included in the **internal** group. The underlying IGP protocol is IS-IS on both routers, and relevant interfaces are configured to run IS-IS.

Note that in the configuration on both routers, the router ID is manually configured to avoid duplicate router ID problems, and the **next-hop-self** statement is included to avoid any BGP next-hop reachability problems.

Verify Advertised BGP Routes

Purpose

You can determine if a particular route that you have configured is being advertised to a neighbor.

Action

To verify the routing information as it has been prepared for advertisement to the specified BGP neighbor, enter the following Junos OS CLI operational mode command:

```
user@host> show route advertising-protocol bgp neighbor-address
```

Sample Output

```
user@R2> show route advertising-protocol bgp 10.0.0.4\
```

```
inet.0: 20 destinations, 22 routes (20 active, 0 holddown, 0 hidden)
  Prefix                Nexthop          MED      Lclpref    AS path
* 100.100.1.0/24        Self             5         200        65001 I
* 100.100.2.0/24        Self             5         100        65001 I
* 100.100.3.0/24        Self             100       100        65001 I
* 100.100.4.0/24        Self             100       100        65001 I
```

Meaning

The sample output shows the BGP routes advertised from **R2** to its neighbor, **10.0.0.4 (R4)**. Out of 22 total routes in the **inet.0** routing table, 20 are active destinations. No routes are **hidden** or in the **hold-down** state. Routes reside in the **hold-down** state prior to being declared active, and routes rejected by a routing policy can be placed into the **hidden** state. The information displayed reflects the routes that the routing table exported to the BGP routing protocol.

Verify That a Particular BGP Route Is Received on Your Router

Purpose

Display the routing information as it is received through a particular BGP neighbor and advertised by the local router to the neighbor.

Action

To verify that a particular BGP route is received on your router, enter the following Junos OS CLI operational mode command:


```
user@host> show route receive-protocol bgp neighbor-address
```

Sample Output

```
user@R6> show route receive-protocol bgp 10.0.0.2
```

```
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
  Prefix                Nexthop                MED      Lclpref    AS path
* 100.100.1.0/24        10.0.0.2              5    200    65001 I
* 100.100.2.0/24        10.0.0.2              5    100    65001 I
  100.100.3.0/24        10.0.0.2              100           65001 I
  100.100.4.0/24        10.0.0.2              100           65001 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

user@R6> show route receive-protocol bgp 10.0.0.4
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
  Prefix                Nexthop                MED      Lclpref    AS path
* 100.100.3.0/24        10.0.0.4              100           65001 I
* 100.100.4.0/24        10.0.0.4              100           65001 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Meaning

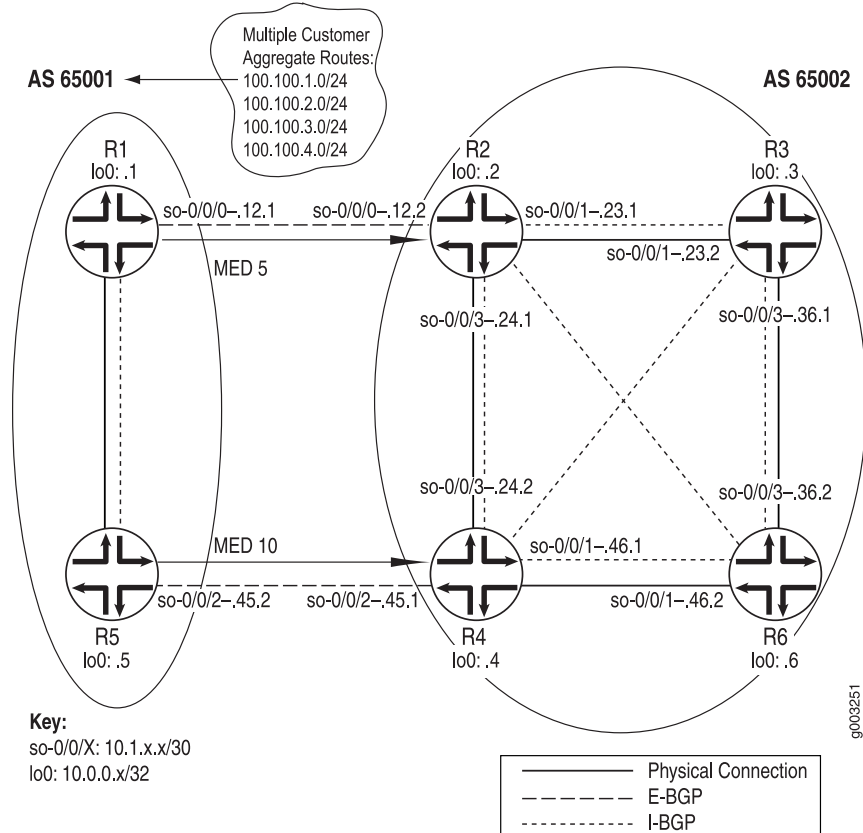
The sample output shows four BGP routes from **R2** and two from **R4**. Of the four routes from **R2**, only two are active in the routing table, as indicated by the asterisk (*), while both routes received from **R4** are active in the routing table. All BGP routes came through AS 65001.

Examine BGP Routes and Route Selection

Purpose

You can examine the BGP path selection process to determine the single, active path when BGP receives multiple routes to the same destination prefix.

Figure 178: BGP Network Topology



The network in [Figure 178 on page 2419](#) shows that **R1** and **R5** announce the same aggregate routes to **R2** and **R4**, which results in **R2** and **R4** receiving two routes to the same destination prefix. The route selection process on **R2** and **R4** determines which of the two BGP routes received is active and advertised to the other internal routers (**R3** and **R6**).

Before the router installs a BGP route, it must make sure that the BGP **next-hop** attribute is reachable. If the BGP next hop cannot be resolved, the route is not installed. When a BGP route is installed in the routing table, it must go through a path selection process if multiple routes exist to the same destination prefix. The BGP path selection process proceeds in the following order:

1. Route preference in the routing table is compared. For example, if an OSPF and a BGP route exist for a particular destination, the OSPF route is selected as the active route because the OSPF route has a default preference of 110, while the BGP route has a default preference of 170.
2. Routes are compared for local preference. The route with the highest local preference is preferred. For example, see [“Examine the Local Preference Selection” on page 2420](#).
3. The AS path attribute is evaluated. The shorter AS path is preferred.
4. The origin code is evaluated. The lowest origin code is preferred (**I (IGP)** < **E (EGP)** < **? (Incomplete)**).

5. The MED value is evaluated. By default, if any of the routes are advertised from the same neighboring AS, the lowest MED value is preferred. The absence of a MED value is interpreted as a MED of 0. For an example, see [“Examine the Multiple Exit Discriminator Route Selection” on page 2422](#).
6. The route is evaluated as to whether it is learned through EBGP or IBGP. EBGP learned routes are preferred to IBGP learned routes. For an example, see [“Examine the EBGP over IBGP Selection” on page 2423](#)
7. If the route is learned from IBGP, the route with the lowest IGP cost is preferred. For an example, see [“Examine the IGP Cost Selection” on page 2425](#). The physical next hop to the IBGP peer is installed according to the following three rules:
 - a. After BGP examines the **inet.0** and **inet.3** routing tables, the physical next hop of the route with the lowest preference is used.
 - b. If the preference values in the **inet.0** and the **inet.3** routing tables are a tie, the physical next hop of the route in the **inet.3** routing table is used.
 - c. When a preference tie exists in the same routing table, the physical next hop of the route with more paths is installed.
8. The route reflection cluster list attribute is evaluated. The shortest length cluster list is preferred. Routes without a cluster list are considered to have a cluster list length of 0.
9. The router ID is evaluated. The route from the peer with the lowest router ID is preferred (usually the loopback address).
10. The peer address value is examined. The peer with the lowest peer IP address is preferred.

To determine the single, active path when BGP receives multiple routes to the same destination prefix, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```

The following steps illustrate the inactive reason displayed when BGP receives multiple routes to the same destination prefix and one route is selected as the single, active path:

1. [Examine the Local Preference Selection | 2420](#)
2. [Examine the Multiple Exit Discriminator Route Selection | 2422](#)
3. [Examine the EBGP over IBGP Selection | 2423](#)
4. [Examine the IGP Cost Selection | 2425](#)

Examine the Local Preference Selection

Purpose

To examine a route to determine if local preference is the selection criteria for the single, active path.

Action

To examine a route to determine if local preference is the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```

Sample Output

```
user@R4> show route 100.100.1.0 detail
```

```
inet.0: 20 destinations, 24 routes (20 active, 0 holddown, 0 hidden)
100.100.1.0/24 (2 entries, 1 announced)
    *BGP      Preference: 170/-201
        Source: 10.0.0.2
        Next hop: 10.1.24.1 via so-0/0/3.0, selected
        Protocol next hop: 10.0.0.2 Indirect next hop: 8644000 277
        State: <Active Int Ext>
        Local AS: 65002 Peer AS: 65002
        Age: 2:22:34      Metric: 5      Metric2: 10
        Task: BGP_65002.10.0.0.2+179
        Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
        AS path: 65001 I
        Localpref: 200
        Router ID: 10.0.0.2
    BGP      Preference: 170/-101
        Source: 10.1.45.2
        Next hop: 10.1.45.2 via so-0/0/2.0, selected
        State: <Ext>
        Inactive reason: Local Preference
        Local AS: 65002 Peer AS: 65001
        Age: 2w0d 1:28:31      Metric: 10
        Task: BGP_65001.10.1.45.2+179
        AS path: 65001 I
        Localpref: 100
        Router ID: 10.0.0.5
```

Meaning

The sample output shows that **R4** received two instances of the **100.100.1.0** route: one from **10.0.0.2 (R2)** and one from **10.1.45.2 (R5)**. **R4** selected the path from **R2** as its active path, as indicated by the asterisk

(*). The selection is based on the local preference value contained in the **Localpref** field. The path with the *highest* local preference is preferred. In the example, the path with the higher local preference value is the path from **R2**, 200.

The reason that the route from **R5** is not selected is in the **Inactive reason** field, in this case, **Local Preference**.

Note that the two paths are from the same neighboring network: AS 65001.

Examine the Multiple Exit Discriminator Route Selection

Purpose

To examine a route to determine if the MED is the selection criteria for the single, active path.

Action

To examine a route to determine if the MED is the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```

Sample Output

```
user@R4> show route 100.100.2.0 detail
```

```
inet.0: 20 destinations, 24 routes (20 active, 0 holddown, 0 hidden)
100.100.2.0/24 (2 entries, 1 announced)
    *BGP      Preference: 170/-101
        Source: 10.0.0.2
        Next hop: 10.1.24.1 via so-0/0/3.0, selected
        Protocol next hop: 10.0.0.2 Indirect next hop: 8644000 277
        State: <Active Int Ext>
        Local AS: 65002 Peer AS: 65002
        Age: 2:32:01      Metric: 5      Metric2: 10
        Task: BGP_65002.10.0.0.2+179
        Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
        AS path: 65001 I
        Localpref: 100
        Router ID: 10.0.0.2
    BGP      Preference: 170/-101
        Source: 10.1.45.2
        Next hop: 10.1.45.2 via so-0/0/2.0, selected
```



```

State: <NotBest Ext>
  Inactive reason: Not Best in its group
Local AS: 65002 Peer AS: 65001
Age: 2w0d 1:37:58      Metric: 10
Task: BGP_65001.10.1.45.2+179
  AS path: 65001 I
Localpref: 100
Router ID: 10.0.0.5

```

Meaning

The sample output shows that **R4** received two instances of the **100.100.2.0** route: one from **10.0.0.2 (R2)**, and one from **10.1.45.2 (R5)**. **R4** selected the path from **R2** as its active route, as indicated by the asterisk (*). The selection is based on the MED value contained in the **Metric:** field. The path with the lowest MED value is preferred. In the example, the path with the lowest MED value (5) is the path from **R2**. Note that the two paths are from the same neighboring network: AS 65001.

The reason that the inactive path is not selected is displayed in the **Inactive reason:** field, in this case, **Not Best in its group**. The wording is used because the Junos OS uses the process of deterministic MED selection, by default.

Examine the EBGW over IBGP Selection

Purpose

To examine a route to determine if EBGW is selected over IBGP as the selection criteria for the single, active path.

Action

To examine a route to determine if EBGW is selected over IBGP as the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```

Sample Output

```
user@R4> show route 100.100.3.0 detail
```

```

inet.0: 20 destinations, 24 routes (20 active, 0 holddown, 0 hidden)
100.100.3.0/24 (2 entries, 1 announced)

```



```

*BGP      Preference: 170/-101
          Source: 10.1.45.2
          Next hop: 10.1.45.2 via so-0/0/2.0, selected
          State: <Active Ext>
          Local AS: 65002 Peer AS: 65001
          Age: 5d 0:31:25
          Task: BGP_65001.10.1.45.2+179
          Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
          AS path: 65001 I
          Localpref: 100
          Router ID: 10.0.0.5
BGP      Preference: 170/-101
          Source: 10.0.0.2
          Next hop: 10.1.24.1 via so-0/0/3.0, selected
          Protocol next hop: 10.0.0.2 Indirect next hop: 8644000 277
          State: <NotBest Int Ext>
          Inactive reason: Interior > Exterior > Exterior via Interior
          Local AS: 65002 Peer AS: 65002
          Age: 2:48:18      Metric2: 10
          Task: BGP_65002.10.0.0.2+179
          AS path: 65001 I
          Localpref: 100
          Router ID: 10.0.0.2

```

Meaning

The sample output shows that **R4** received two instances of the **100.100.3.0** route: one from **10.1.45.2 (R5)** and one from **10.0.0.2 (R2)**. **R4** selected the path from **R5** as its active path, as indicated by the asterisk (*). The selection is based on a preference for routes learned from an EBGp peer over routes learned from an IBGP. **R5** is an EBGp peer.

You can determine if a path is received from an EBGp or IBGP peer by examining the **Local As** and **Peer As** fields. For example, the route from **R5** shows the local AS is 65002 and the peer AS is 65001, indicating that the route is received from an EBGp peer. The route from **R2** shows that both the local and peer AS is 65002, indicating that it is received from an IBGP peer.

The reason that the inactive path is not selected is displayed in the **Inactive reason** field, in this case, **Interior > Exterior > Exterior via Interior**. The wording of this reason shows the order of preferences applied when the same route is received from two routers. The route received from a strictly internal source (IGP) is preferred first, the route received from an external source (EBGP) is preferred next, and any route which comes from an external source and is received internally (IBGP) is preferred last.

Examine the IGP Cost Selection

Purpose

To examine a route to determine if EBGp is selected over IBGP as the selection criteria for the single, active path.

Action

To examine a route to determine if EBGp is selected over IBGP as the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```

Sample Output

```
user@R6> show route 100.100.4.0 detail
```

```
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
100.100.4.0/24 (2 entries, 1 announced)
  *BGP      Preference: 170/-101
    Source: 10.0.0.4
      Next hop: 10.1.46.1 via so-0/0/1.0, selected
      Protocol next hop: 10.0.0.4 Indirect next hop: 864c000 276
      State: <Active Int Ext>
      Local AS: 65002 Peer AS: 65002
      Age: 2:16:11      Metric2: 10
      Task: BGP_65002.10.0.0.4+4120
      Announcement bits (2): 0-KRT 4-Resolve inet.0
    AS path: 65001 I
      Localpref: 100
      Router ID: 10.0.0.4
  BGP      Preference: 170/-101
    Source: 10.0.0.2
      Next hop: 10.1.46.1 via so-0/0/1.0, selected
      Next hop: 10.1.36.1 via so-0/0/3.0
      Protocol next hop: 10.0.0.2 Indirect next hop: 864c0b0 278
      State: <NotBest Int Ext>
    Inactive reason: IGP metric
      Local AS: 65002 Peer AS: 65002
      Age: 2:16:03      Metric2: 20
      Task: BGP_65002.10.0.0.2+179
```



```

AS path: 65001 |
Localpref: 100
Router ID: 10.0.0.2

```

Meaning

The sample output shows that **R6** received two instances of the **100.100.4.0** route: one from **10.0.0.4 (R4)** and one from **10.0.0.2 (R2)**. **R6** selected the path from **R4** as its active route, as indicated by the asterisk (*). The selection is based on the IGP metric, displayed in the **Metric2** field. The route with the lowest IGP metric is preferred. In the example, the path with the lowest IGP metric value is the path from **R4**, with an IGP metric value of 10, while the path from **R2** has an IGP metric of 20. Note that the two paths are from the same neighboring network: AS 65001.

The reason that the inactive path was not selected is displayed in the **Inactive reason** field, in this case, **IGP metric**.

Examine the Local Preference Selection

Purpose

To examine a route to determine if local preference is the selection criteria for the single, active path.

Action

To examine a route to determine if local preference is the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```

Sample Output

```
user@R4> show route 100.100.1.0 detail
```

```

inet.0: 20 destinations, 24 routes (20 active, 0 holddown, 0 hidden)
100.100.1.0/24 (2 entries, 1 announced)
    *BGP      Preference: 170/-201
              Source: 10.0.0.2
              Next hop: 10.1.24.1 via so-0/0/3.0, selected
              Protocol next hop: 10.0.0.2 Indirect next hop: 8644000 277
              State: <Active Int Ext>

```



```

Local AS: 65002 Peer AS: 65002
Age: 2:22:34      Metric: 5      Metric2: 10
Task: BGP_65002.10.0.0.2+179
Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
AS path: 65001 I
Localpref: 200
Router ID: 10.0.0.2
BGP Preference: 170/-101
Source: 10.1.45.2
Next hop: 10.1.45.2 via so-0/0/2.0, selected
State: <Ext>
Inactive reason: Local Preference
Local AS: 65002 Peer AS: 65001
Age: 2w0d 1:28:31      Metric: 10
Task: BGP_65001.10.1.45.2+179
AS path: 65001 I
Localpref: 100
Router ID: 10.0.0.5

```

Meaning

The sample output shows that **R4** received two instances of the **100.100.1.0** route: one from **10.0.0.2 (R2)** and one from **10.1.45.2 (R5)**. **R4** selected the path from **R2** as its active path, as indicated by the asterisk (*). The selection is based on the local preference value contained in the **Localpref** field. The path with the *highest* local preference is preferred. In the example, the path with the higher local preference value is the path from **R2**, 200.

The reason that the route from **R5** is not selected is in the **Inactive reason** field, in this case, **Local Preference**.

Note that the two paths are from the same neighboring network: AS 65001.

Examine the Multiple Exit Discriminator Route Selection

Purpose

To examine a route to determine if the MED is the selection criteria for the single, active path.

Action

To examine a route to determine if the MED is the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:


```
user@host> show route destination-prefix < detail >
```

Sample Output

```
user@R4> show route 100.100.2.0 detail
```

```
inet.0: 20 destinations, 24 routes (20 active, 0 holddown, 0 hidden)
100.100.2.0/24 (2 entries, 1 announced)
    *BGP      Preference: 170/-101
        Source: 10.0.0.2
        Next hop: 10.1.24.1 via so-0/0/3.0, selected
        Protocol next hop: 10.0.0.2 Indirect next hop: 8644000 277
        State: <Active Int Ext>
        Local AS: 65002 Peer AS: 65002
        Age: 2:32:01      Metric: 5      Metric2: 10
        Task: BGP_65002.10.0.0.2+179
        Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
        AS path: 65001 I
        Localpref: 100
        Router ID: 10.0.0.2
    BGP      Preference: 170/-101
        Source: 10.1.45.2
        Next hop: 10.1.45.2 via so-0/0/2.0, selected
        State: <NotBest Ext>
        Inactive reason: Not Best in its group
        Local AS: 65002 Peer AS: 65001
        Age: 2w0d 1:37:58      Metric: 10
        Task: BGP_65001.10.1.45.2+179
        AS path: 65001 I
        Localpref: 100
        Router ID: 10.0.0.5
```

Meaning

The sample output shows that **R4** received two instances of the **100.100.2.0** route: one from **10.0.0.2 (R2)**, and one from **10.1.45.2 (R5)**. **R4** selected the path from **R2** as its active route, as indicated by the asterisk (*). The selection is based on the MED value contained in the **Metric:** field. The path with the lowest MED value is preferred. In the example, the path with the lowest MED value (5) is the path from **R2**. Note that the two paths are from the same neighboring network: AS 65001.

The reason that the inactive path is not selected is displayed in the **Inactive reason:** field, in this case, **Not Best in its group**. The wording is used because the Junos OS uses the process of deterministic MED selection, by default.

Examine the EBGp over IBGP Selection

Purpose

To examine a route to determine if EBGp is selected over IBGP as the selection criteria for the single, active path.

Action

To examine a route to determine if EBGp is selected over IBGP as the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```

Sample Output

```
user@R4> show route 100.100.3.0 detail
```

```
inet.0: 20 destinations, 24 routes (20 active, 0 holddown, 0 hidden)
100.100.3.0/24 (2 entries, 1 announced)
    *BGP      Preference: 170/-101
        Source: 10.1.45.2
        Next hop: 10.1.45.2 via so-0/0/2.0, selected
        State: <Active Ext>
        Local AS: 65002 Peer AS: 65001
        Age: 5d 0:31:25
        Task: BGP_65001.10.1.45.2+179
        Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
        AS path: 65001 I
        Localpref: 100
        Router ID: 10.0.0.5
    BGP      Preference: 170/-101
        Source: 10.0.0.2
        Next hop: 10.1.24.1 via so-0/0/3.0, selected
        Protocol next hop: 10.0.0.2 Indirect next hop: 8644000 277
        State: <NotBest Int Ext>
```


Inactive reason: Interior > Exterior > Exterior via Interior

Local AS: 65002 Peer AS: 65002

Age: 2:48:18 Metric2: 10

Task: BGP_65002.10.0.0.2+179

AS path: 65001 I

Localpref: 100

Router ID: 10.0.0.2

Meaning

The sample output shows that **R4** received two instances of the **100.100.3.0** route: one from **10.1.45.2 (R5)** and one from **10.0.0.2 (R2)**. **R4** selected the path from **R5** as its active path, as indicated by the asterisk (*). The selection is based on a preference for routes learned from an EBGp peer over routes learned from an IBGP. **R5** is an EBGp peer.

You can determine if a path is received from an EBGp or IBGP peer by examining the **Local As** and **Peer As** fields. For example, the route from **R5** shows the local AS is 65002 and the peer AS is 65001, indicating that the route is received from an EBGp peer. The route from **R2** shows that both the local and peer AS is 65002, indicating that it is received from an IBGP peer.

The reason that the inactive path is not selected is displayed in the **Inactive reason** field, in this case, **Interior > Exterior > Exterior via Interior**. The wording of this reason shows the order of preferences applied when the same route is received from two routers. The route received from a strictly internal source (IGP) is preferred first, the route received from an external source (EBGP) is preferred next, and any route which comes from an external source and is received internally (IBGP) is preferred last.

Examine the IGP Cost Selection

Purpose

To examine a route to determine if EBGp is selected over IBGP as the selection criteria for the single, active path.

Action

To examine a route to determine if EBGp is selected over IBGP as the selection criteria for the single, active path, enter the following Junos OS CLI operational mode command:

```
user@host> show route destination-prefix < detail >
```


Sample Output

user@R6> show route 100.100.4.0 detail

```
inet.0: 18 destinations, 20 routes (18 active, 0 holddown, 0 hidden)
100.100.4.0/24 (2 entries, 1 announced)
    *BGP      Preference: 170/-101
              Source: 10.0.0.4
              Next hop: 10.1.46.1 via so-0/0/1.0, selected
              Protocol next hop: 10.0.0.4 Indirect next hop: 864c000 276
              State: <Active Int Ext>
              Local AS: 65002 Peer AS: 65002
              Age: 2:16:11      Metric2: 10
              Task: BGP_65002.10.0.0.4+4120
              Announcement bits (2): 0-KRT 4-Resolve inet.0
              AS path: 65001 I
              Localpref: 100
              Router ID: 10.0.0.4
    BGP      Preference: 170/-101
              Source: 10.0.0.2
              Next hop: 10.1.46.1 via so-0/0/1.0, selected
              Next hop: 10.1.36.1 via so-0/0/3.0
              Protocol next hop: 10.0.0.2 Indirect next hop: 864c0b0 278
              State: <NotBest Int Ext>
              Inactive reason: IGP metric
              Local AS: 65002 Peer AS: 65002
              Age: 2:16:03      Metric2: 20
              Task: BGP_65002.10.0.0.2+179
              AS path: 65001 I
              Localpref: 100
              Router ID: 10.0.0.2
```

Meaning

The sample output shows that **R6** received two instances of the **100.100.4.0** route: one from **10.0.0.4 (R4)** and one from **10.0.0.2 (R2)**. **R6** selected the path from **R4** as its active route, as indicated by the asterisk (*). The selection is based on the IGP metric, displayed in the **Metric2** field. The route with the lowest IGP metric is preferred. In the example, the path with the lowest IGP metric value is the path from **R4**, with an IGP metric value of 10, while the path from **R2** has an IGP metric of 20. Note that the two paths are from the same neighboring network: AS 65001.

The reason that the inactive path was not selected is displayed in the **Inactive reason** field, in this case, **IGP metric**.

Examine Routes in the Forwarding Table

Purpose

When you run into problems, such as connectivity problems, you may need to examine routes in the forwarding table to verify that the routing protocol process has relayed the correct information into the forwarding table.

Action

To display the set of routes installed in the forwarding table, enter the following Junos OS CLI operational mode command:

```
user@host> show route forwarding-table
```

Sample Output

```
user@R2> show route forwarding-table
```

```
Routing table: inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
default              perm    0                               rjct   10    1
10.0.0.2/32          intf    0 10.0.0.2              locl   256    1
10.0.0.3/32          user    1 10.1.23.0              ucst   282    4 so-0/0/1.0
10.0.0.4/32          user    1 10.1.24.0              ucst   290    7 so-0/0/3.0
10.0.0.6/32          user    1 10.1.24.0              ucst   290    7 so-0/0/3.0
10.1.12.0/30         intf    1 ff.3.0.21             ucst   278    6 so-0/0/0.0
10.1.12.0/32         dest    0 10.1.12.0              recv   280    1 so-0/0/0.0
10.1.12.2/32         intf    0 10.1.12.2              locl   277    1
10.1.12.3/32         dest    0 10.1.12.3              bcst   279    1 so-0/0/0.0
10.1.23.0/30         intf    0 ff.3.0.21             ucst   282    4 so-0/0/1.0
10.1.23.0/32         dest    0 10.1.23.0              recv   284    1 so-0/0/1.0
10.1.23.1/32         intf    0 10.1.23.1              locl   281    1
10.1.23.3/32         dest    0 10.1.23.3              bcst   283    1 so-0/0/1.0
10.1.24.0/30         intf    0 ff.3.0.21             ucst   290    7 so-0/0/3.0
10.1.24.0/32         dest    0 10.1.24.0              recv   292    1 so-0/0/3.0
10.1.24.1/32         intf    0 10.1.24.1              locl   289    1
10.1.24.3/32         dest    0 10.1.24.3              bcst   291    1 so-0/0/3.0
10.1.36.0/30         user    0 10.1.23.0              ucst   282    4 so-0/0/1.0
10.1.46.0/30         user    0 10.1.24.0              ucst   290    7 so-0/0/3.0
100.100.1.0/24       user    0 10.1.12.0              ucst   278    6 so-0/0/0.0
100.100.2.0/24       user    0 10.1.12.0              ucst   278    6 so-0/0/0.0
```



```

100.100.3.0/24      user      0 10.1.12.0      ucst    278      6 so-0/0/0.0
100.100.4.0/24      user      0 10.1.12.0      ucst    278      6 so-0/0/0.0
[...Output truncated...]

```

Meaning

The sample output shows the network-layer prefixes and their next hops installed in the forwarding table. The output includes the same next-hop information as in the **show route detail** command (the next-hop address and interface name). Additional information includes the destination type, the next-hop type, the number of references to this next hop, and an index into an internal next-hop database. (The internal database contains additional information used by the Packet Forwarding Engine to ensure proper encapsulation of packets sent out an interface. This database is not accessible to the user.

For detailed information about the meanings of the various flags and types fields, see the [Routing Policies, Firewall Filters, and Traffic Policers User Guide](#).

Ping the Egress Router

Purpose

Ping the egress router to confirm that communication over the network is operational.

Action

To ping the egress router, enter the following Junos OS CLI operational mode command:

```
user@host> ping ip-address-interface
```

Sample Output

```

[edit protocols mpls]
user@R1# run ping 10.1.56.1
PING 10.1.56.1 (10.1.56.1): 56 data bytes
64 bytes from 10.1.56.1: icmp_seq=0 ttl=255 time=0.837 ms
64 bytes from 10.1.56.1: icmp_seq=1 ttl=255 time=0.792 ms
64 bytes from 10.1.56.1: icmp_seq=2 ttl=255 time=0.856 ms
^C
--- 10.1.56.1 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.792/0.828/0.856/0.027 ms

```

Meaning

The sample output confirms that communication between router **R1** and the IP address of the relevant interface on router **R5 (10.1.56.1)** is operational.

View the RSVP Log File on Transit Routers

Purpose

Transit router messages that appear in the RSVP log file can help you analyze the problem with an RSVP session. You may need to issue the **clear rsvp session** and **clear log filename** commands to ensure that your records are current. However, if your network is large with many RSVP sessions, this may not be advisable because it may take a while for all sessions to reestablish. However, the **clear rsvp session** command has various options you can include to minimize the effect on your network. For more information about the **clear rsvp session** command, see the *Junos Routing Protocols and Policies Command Reference*.

Action

To view the RSVP log file, enter the following Junos OS CLI operational mode commands:

```
user@host> clear rsvp session (Optional)
user@host> clear log filename (Optional)
user@host> show log filename
```

Sample Output 1

```
user@R3> clear rsvp session

user@R3> clear log rsvp-log

user@R3> show log rsvp-log
Jul 21 16:51:23 R3 clear-log[30656]: logfile cleared
Jul 21 16:51:24   RSVP recv Path 10.0.0.1->10.0.0.5 Len=208 so-0/0/2.0
Jul 21 16:51:24   Session7 Len 16 10.0.0.5(port/tunnel ID 11956) Proto 0
Jul 21 16:51:24   Hop      Len 12 10.1.13.1/0x086cd198
Jul 21 16:51:24   Time     Len  8 30000 ms
Jul 21 16:51:24   SrcRoute Len 20 10.1.13.2 S 10.1.36.2 S
Jul 21 16:51:24   LabelRequest Len  8 EtherType 0x800
Jul 21 16:51:24   Properties Len 12 Primary path
Jul 21 16:51:24   SessionAttribute Len 16 Prio (7,0) flag 0x0 "R1-to-R5"
Jul 21 16:51:24   Sender7   Len 12 10.0.0.1(port/lsp ID 32)
Jul 21 16:51:24   Tspec     Len 36 rate 0bps size 0bps peak Infbps m 20 M 1500
```



```

Jul 21 16:51:24   ADspec   Len 48 MTU 1500
Jul 21 16:51:24   RecRoute Len 12  10.1.13.1
Jul 21 16:51:24 RSVP send Path 10.0.0.1->10.0.0.5 Len=208 so-0/0/3.0
Jul 21 16:51:24   Session7 Len 16 10.0.0.5(port/tunnel ID 11956) Proto 0
Jul 21 16:51:24   Hop       Len 12 10.1.36.1/0x08680264
Jul 21 16:51:24   Time      Len  8 30000 ms
Jul 21 16:51:24   SrcRoute Len 12  10.1.36.2 S
Jul 21 16:51:24   LabelRequest Len  8 EtherType 0x800
Jul 21 16:51:24   Properties Len 12 Primary path
Jul 21 16:51:24   SessionAttribute Len 16 Prio (7,0) flag 0x0 "R1-to-R5"
Jul 21 16:51:24   Sender7   Len 12 10.0.0.1(port/lsp ID 32)
Jul 21 16:51:24   Tspec     Len 36 rate 0bps size 0bps peak Infbps m 20 M 1500
Jul 21 16:51:24   ADspec    Len 48 MTU 1500
Jul 21 16:51:24   RecRoute Len 20 10.1.36.1 10.1.13.1

```

Sample Output 2

```

user@R6> clear rsvp session

user@R6> clear log rsvp-log

user@R6> show log rsvp-log
Jul 21 17:01:21 R6 clear-log[41496]: logfile cleared
Jul 21 17:01:23 RSVP recv Path 10.0.0.1->10.0.0.5 Len=208 so-0/0/3.0
Jul 21 17:01:23   Session7 Len 16 10.0.0.5(port/tunnel ID 11956) Proto 0
Jul 21 17:01:23   Hop       Len 12 10.1.36.1/0x08680264
Jul 21 17:01:23   Time      Len  8 30000 ms
Jul 21 17:01:23   SrcRoute Len 12  10.1.36.2 S
Jul 21 17:01:23   LabelRequest Len  8 EtherType 0x800
Jul 21 17:01:23   Properties Len 12 Primary path
Jul 21 17:01:23   SessionAttribute Len 16 Prio (7,0) flag 0x0 "R1-to-R5"
Jul 21 17:01:23   Sender7   Len 12 10.0.0.1(port/lsp ID 32)
Jul 21 17:01:23   Tspec     Len 36 rate 0bps size 0bps peak Infbps m 20 M 1500
Jul 21 17:01:23   ADspec    Len 48 MTU 1500
Jul 21 17:01:23   RecRoute Len 20 10.1.36.1 10.1.13.1
Jul 21 17:01:23 RSVP send Path 10.0.0.1->10.0.0.5 Len=204 so-0/0/0.0
Jul 21 17:01:23   Session7 Len 16 10.0.0.5(port/tunnel ID 11956) Proto 0
Jul 21 17:01:23   Hop       Len 12 10.1.56.2/0x086f9000
Jul 21 17:01:23   Time      Len  8 30000 ms

```



```

Jul 21 17:01:23 LabelRequest Len 8 EtherType 0x800
Jul 21 17:01:23 Properties Len 12 Primary path
Jul 21 17:01:23 SessionAttribute Len 16 Prio (7,0) flag 0x0 "R1-to-R5"
Jul 21 17:01:23 Sender7 Len 12 10.0.0.1(port/lsp ID 32)
Jul 21 17:01:23 Tspec Len 36 rate 0bps size 0bps peak Infbps m 20 M 1500
Jul 21 17:01:23 ADspec Len 48 MTU 1500
Jul 21 17:01:23 RecRoute Len 28 10.1.56.2 10.1.36.1 10.1.13.1

```

Meaning

Sample Output 1 from transit router **R3** shows that **R3 (so-0/0/2.0)** correctly received a Path request message (**recv Path**) from **R1**, and correctly sent the Path message (**send Path**) through interface **so-0/0/3.0** to **R6**. The route record object (**RecRoute**) indicates the list of addresses this Path message transited, in this case, **10.1.36.1** and **10.1.13.1**.

Sample Output 2 from transit router **R6** shows that **R6 (so-0/0/3.0)** correctly received a Path request message (**recv Path**) from **R3**, and correctly sent the Path message (**send Path**) through interface **so-0/0/0** to **R5**. The route record object (**RecRoute**) indicates the list of addresses this Path message transited, in this case, **10.1.56.2**, **10.1.36.1**, and **10.1.13.1**.

With the information above, the focus shifts to egress router **R5** as the source of the problem, with indications that **R5** ignored the RSVP message.

Check the RSVP Log File on the Egress Router

Purpose

After placing an RSVP tracing configuration on router **R5** similar to that on routers **R3** and **R6**, display the RSVP log file for useful information about the problem on router **R5**.

Action

To check the RSVP log file, enter the following Junos OS CLI operational mode command:

```
user@host> show log rsvp-log
```

Sample Output

```

user@R5> show log rsvp-log
Jul 21 10:53:16 R5 clear-log[40071]: logfile cleared

```



```

Jul 21 11:02:37 trace_on: Tracing to "/var/log/rsvp-log" started
Jul 21 11:03:07 RSVP error, send to DISABLED interface ? Hello New
10.1.56.1->10.1.56.2 Len=8  so-0/0/0.0

```

Meaning

The sample output shows that **R5** did not receive the Path message because of a disabled interface, **so-0/0/0.0**.

Determine and Correct the Problem on the Egress Router

Problem

Description: Check the configuration of interface **so-0/0/0.0** on egress router **R5** to determine the reason it was disabled.

Solution

To determine the problem on **R5**, enter the following Junos OS CLI commands:

```

user@R5> show configuration protocols rsvp
user@R5> edit
[edit protocols rsvp]
user@R5# rename interface so-0/0/3 to interface so-0/0/0
user@R5# show
user@R5# commit
user@R5# run show rsvp session ingress detail

```

Sample Output 1

```

user@R5> show configuration protocols rsvp
traceoptions {
  file rsvp-log;
  flag error detail;
  flag path detail;
  flag pathtear detail;
}
interface so-0/0/3.0;      <<< so-0/0/3 incorrectly included
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {

```



```

    disable;
}

```

Sample Output 2

```

[edit protocols rsvp]
user@R5# rename interface so-0/0/3 to interface so-0/0/0

[edit protocols rsvp]
user@R5# show
traceoptions {
    file rsvp-log;
    flag packets detail;
    flag error detail;
}
interface so-0/0/0.0;
interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}

[edit protocols rsvp]
user@R5# commit
commit complete

```

Sample Output 3

```

[edit protocols mpls]
user@R5# run show rsvp session ingress detail
Ingress RSVP: 1 sessions

```

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.0.0.1	10.0.0.5	Up	1	1 FF	-	103104	R5-to-R1

```

Total 1 displayed, Up 1, Down 0
Egress RSVP: 1 sessions

```

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.0.0.5	10.0.0.1	Up	0	1 FF	3	-	R1-to-R5

```

Total 1 displayed, Up 1, Down 0

```



```
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning

Sample Output 1 from egress router **R5** shows three interfaces configured at the `[edit protocols rsvp]` hierarchy level, none of which is **so-0/0/0.0**. On examination of the network topology, it is apparent that the **so-0/0/0.0** interface was configured incorrectly as **so-0/0/3.0**.

Sample Output 2 shows the correct configuration of interfaces at the `[edit protocols rsvp]` hierarchy level, and the **rename** command issued to correct the configuration error.

Sample Output 3 shows that the RSVP-signaled LSP (**R1-to-R5**) is correctly established after the changes to the RSVP configuration are committed.

Check the Routing CPU Memory Usage

Purpose

Software processes on the router can consume a considerable amount of CPU and memory. The routing protocol process (rpd) can consume enormous amounts of memory to store information needed for the operation of routing and related protocols, such as Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (ISIS), Resource Reservation Protocol (RSVP), Label Distribution Protocol (LDP), and Multiprotocol Label Switching (MPLS).

To verify the traffic passing through the router and check memory utilization, follow these steps:

1. [Check Overall CPU and Memory Usage | 2439](#)
2. [Check Routing Protocol Process \(rpd\) Memory Usage | 2442](#)
3. [Display Tasks | 2446](#)

Check Overall CPU and Memory Usage

Purpose

You can display exhaustive system process information about software processes that are running on the router and have controlling terminals. This command is equivalent to the UNIX **top** command. However, the UNIX **top** command shows real-time memory usage, with the memory values constantly changing, while the **show system processes extensive** command provides a snapshot of memory usage in a given moment.

Action

To check overall CPU and memory usage, enter the following Junos OS command-line interface (CLI) command:

```
user@host> show system processes extensive
```

Sample Output

```
user@R1> show system processes extensive
```

```
last pid: 5251; load averages: 0.00, 0.00, 0.00 up 4+20:22:16 10:44:41
```

```
58 processes: 1 running, 57 sleeping
```

```
Mem: 57M Active, 54M Inact, 17M Wired, 184K Cache, 35M Buf, 118M Free
```

```
Swap: 512M Total, 512M Free
```

PID	USERNAME	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	CPU	COMMAND
4480	root	2	0	3728K	1908K	select	231:17	2.34%	2.34%	chassisd
4500	root	2	0	1896K	952K	select	0:36	0.00%	0.00%	fud
4505	root	2	0	1380K	736K	select	0:35	0.00%	0.00%	irsd
4481	root	2	0	1864K	872K	select	0:32	0.00%	0.00%	alarmd
4488	root	2	0	8464K	4600K	kqread	0:28	0.00%	0.00%	rpd
4501	root	2	-15	1560K	968K	select	0:21	0.00%	0.00%	ppmd
4510	root	2	0	1372K	812K	select	0:13	0.00%	0.00%	bfdd
5	root	18	0	0K	0K	syncer	0:09	0.00%	0.00%	syncer
4485	root	2	0	3056K	1776K	select	0:07	0.00%	0.00%	snmpd
4499	root	2	0	3688K	1676K	select	0:05	0.00%	0.00%	kmd
4486	root	2	0	3760K	1748K	select	0:05	0.00%	0.00%	mib2d
4493	root	2	0	1872K	928K	select	0:03	0.00%	0.00%	pfed
4507	root	2	0	1984K	1052K	select	0:02	0.00%	0.00%	fsad
4518	root	2	0	3780K	2400K	select	0:02	0.00%	0.00%	dcd
8	root	-18	0	0K	0K	psleep	0:02	0.00%	0.00%	vmuncachedaemo
4	root	-18	0	0K	0K	psleep	0:02	0.00%	0.00%	bufdaemon
4690	root	2	0	0K	0K	peer_s	0:01	0.00%	0.00%	peer proxy
4504	root	2	0	1836K	968K	select	0:01	0.00%	0.00%	dfwd
4477	root	2	0	992K	320K	select	0:01	0.00%	0.00%	watchdog
4354	root	2	0	1116K	604K	select	0:01	0.00%	0.00%	syslogd
4492	root	10	0	1004K	400K	nanslp	0:01	0.00%	0.00%	tnp.snmpd
4446	root	10	0	1108K	616K	nanslp	0:01	0.00%	0.00%	cron
4484	root	2	0	15716K	7468K	select	0:01	0.00%	0.00%	mgd
4494	root	2	15	2936K	2036K	select	0:01	0.00%	0.00%	sampled
5245	remote	2	0	8340K	3472K	select	0:01	0.00%	0.00%	cli
2	root	-18	0	0K	0K	psleep	0:00	0.00%	0.00%	pagedaemon
4512	root	2	0	2840K	1400K	select	0:00	0.00%	0.00%	l2tpd


```

    1 root      10   0   852K   580K wait    0:00  0.00%  0.00% init
5244 root       2   0  1376K   784K select   0:00  0.00%  0.00% telnetd
4509 root      10   0  1060K   528K nanslp   0:00  0.00%  0.00% eccd
4508 root       2   0  2264K  1108K select   0:00  0.00%  0.00% spd
2339 root      10   0   514M 17260K mfsidl   0:00  0.00%  0.00% newfs
4497 root       2   0  2432K  1152K select   0:00  0.00%  0.00% cosd
4490 root       2 -15  2356K  1020K select   0:00  0.00%  0.00% apsd
4496 root       2   0  2428K  1108K select   0:00  0.00%  0.00% rmopd
4491 root       2   0  2436K  1104K select   0:00  0.00%  0.00% vrrpd
4487 root       2   0 15756K  7648K sbwait   0:00  0.00%  0.00% mgd
5246 root       2   0 15776K  8336K select   0:00  0.00%  0.00% mgd
    0 root     -18   0     0K     0K sched    0:00  0.00%  0.00% swapper
5251 root    30  0 21732K  840K RUN    0:00  0.00%  0.00% top
4511 root       2   0  1964K   908K select   0:00  0.00%  0.00% pgmd
4502 root       2   0  1960K   956K select   0:00  0.00%  0.00% lmpd
4495 root       2   0  1884K   876K select   0:00  0.00%  0.00% ilmid
4482 root       2   0  1772K   776K select   0:00  0.00%  0.00% craftd
4503 root      10   0  1040K   492K nanslp   0:00  0.00%  0.00% smartd
    6 root     28   0     0K     0K sleep    0:00  0.00%  0.00% netdaemon
4498 root       2   0  1736K   932K select   0:00  0.00%  0.00% nasd
4506 root       2   0  1348K   672K select   0:00  0.00%  0.00% rtspd
4489 root       2   0  1160K   668K select   0:00  0.00%  0.00% inetd
4478 root       2   0  1108K   608K select   0:00  0.00%  0.00% tnetd
4483 root       2   0  1296K   540K select   0:00  0.00%  0.00% ntpd
4514 root       3   0  1080K   540K ttyin    0:00  0.00%  0.00% getty
4331 root       2   0   416K   232K select   0:00  0.00%  0.00% pccardd
    7 root       2   0     0K     0K pfeacc   0:00  0.00%  0.00% if_pfe_listen
   11 root       2   0     0K     0K picacc   0:00  0.00%  0.00% if_pic_listen
    3 root      18   0     0K     0K psleep   0:00  0.00%  0.00% vmdaemon
    9 root       2   0     0K     0K scs_ho   0:00  0.00%  0.00% scs_housekeepi
   10 root       2   0     0K     0K cb-pol   0:00  0.00%  0.00% cb_poll

```

Meaning

The sample output shows the amount of virtual memory used by the Routing Engine and software processes. For example, 118 MB of physical memory is free and 512 MB of the swap file is free, indicating that the router is not short of memory. The **processes** field shows that most of the 58 processes are in the **sleeping** state, with 1 in the **running** state. The process or command that is running is the **top** command.

The **commands** column lists the processes that are currently running. For example, the chassis process (chassisd) has a process identifier (**PID**) of 4480, with a current priority (**PRI**) of 2. A lower priority number indicates a higher priority.

The processes are listed according to level of activity, with the most active process at the top of the output. For example, the chassis (chassisd) process is consuming the largest amount of CPU resource at 2.34 percent.

The memory field (**Mem**) shows the virtual memory managed by the Routing Engine and used by processes. The value in the memory field is in KB and MB, and is broken down as follows:

- **Active**—Memory that is allocated and actually in use by programs.
- **Inact**—Memory that is either allocated but not recently used or memory that was freed by programs. Inactive memory is still mapped in the address space of one or more processes and, therefore, counts toward the resident set size of those processes.
- **Wired**—Memory that is not eligible to be swapped, and is usually used for Routing Engine memory structures or memory physically locked by a process.
- **Cache**—Memory that is not associated with any program and does not need to be swapped before being reused.
- **Buf**—The size of the memory buffer used to hold data recently called from disk.
- **Free**—Memory that is not associated with any programs. Memory freed by a process can become **Inactive**, **Cache**, or **Free**, depending on the method used by the process to free the memory.

When the system is under memory pressure, the pageout process reuses memory from the free, cache, inactive and, if necessary, active pages.

The **Swap** field shows the total swap space available and how much is unused. In the example, the output shows 512 MB of total swap space and 512 MB of free swap space.

Finally, the memory usage of each process is listed. The **SIZE** field indicates the size of the virtual address space, and the **RES** field indicates the amount of the program in physical memory, which is also known as RSS or Resident Set Size. In the sample output, the chassis (chassisd) process has 3728 KB of virtual address space and 1908 KB of physical memory.

Check Routing Protocol Process (rpd) Memory Usage

Purpose

When you notice a lot of memory usage, you can obtain detailed information about the memory utilization of routing tasks to get an idea of what is going on. The routing process (rpd) is the main task that uses Routing Engine memory.

Action

To check routing process memory usage, enter the following Junos OS CLI operational mode commands:

```
user@host> show route summary
user@host> show task memory detail
```


Sample Output

user@host> **show route summary**

```

Autonomous system number: 209
Router ID: 205.175.0.170
inet.0: 179783 destinations, 898393 routes (179771 active, 146 holddown, 157 hidden)

      Direct:      17 routes,      17 active
      Local:      18 routes,      18 active
      BGP: 896632 routes, 178010 active
      Static:     32 routes,     31 active
      IS-IS:    1694 routes,    1694 active
inet.2: 8766 destinations, 22700 routes (8766 active, 124 holddown, 73 hidden)
      Direct:      17 routes,      17 active
      Local:      18 routes,      18 active
      BGP:  20939 routes,   7006 active
      Static:     32 routes,     31 active
      IS-IS:    1694 routes,    1694 active
inet.3: 1614 destinations, 1719 routes (1614 active, 0 holddown, 0 hidden)
      IS-IS:    1613 routes,    1551 active
      RSVP:     45 routes,     45 active
      LDP:      61 routes,     18 active
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Direct:     1 routes,     1 active
mpls.0: 371 destinations, 371 routes (371 active, 0 holddown, 0 hidden)
      MPLS:       3 routes,       3 active
      RSVP:    303 routes,    303 active
      LDP:     65 routes,     65 active

```

user@R1> **show task memory detail**

```

----- Overall Memory Report -----
Size TP    Allocs  Mallocs  AllocBytes  MaxAllocs  MaxBytes  FreeBytes
  12      8140   186959   2341188    200824     2409888    54972
  16     4061    182     67888     4586      73376     5840
  16 T         -      -         -    393571    6297136         -
  20    688588    51   13772780    713704   14274080   423956
[...Output truncated...]
 8192 P      91      -    745472     195    1597440         -
12288 P       -      -         -      1     12288         -
block      5      -    137200     14    137732     6160
pool     50      -     896    100     1792     3200
alloc     -      8   383744     10   397365     9472
-----

```



```

389169664      578341705  72977920
----- Allocator Memory Report -----
Name           Size Alloc DTP      Alloc      Alloc MaxAlloc  MaxAlloc
                Size      Blocks    Bytes    Blocks    Bytes
patricia_root_struct  8   12      7741    92892     8130    97560
sockaddr_un.i802      8   12         2      24         2      24
sockaddr_un.tag       8   12      371    4452      995    11940
if_addr_entry         8   12         -      -         1      12
gw_entry_list         8   12         1      12         1      12
isis_proto_list       8   12        25     300        30     360
struct krt_scb        12  16         4      64         6      96
ldp_rt_data          12  16        61     976       133    2128
config_list          12  16     2353   37648    2353   37648
TED NodeInfo         12  16      845   13520     907   14512
isis_area_addr  12  16    544  8704  612  9792
isis_nh_list          12  16        237   3792     922   14752
isis_tsi              12  16         17     272        19     304
bgp_use_block         12  16         -      -       112   1792
isis_route_walk_cont  12  16  T         -      -         1      16
bgpg_rtinfo_entry     12  16  T         -      -    393571  6297136
task_floating_socket  16  20         1      20         1      20
[...Output truncated...]
rt_parse_memory      4092 4096  TP         -      -         1    4096
noblock_buffer_blk   4092 4096  TP         5    20480     811  3321856
bgp_buffer           4100 8192   P        91   745472     100  819200
bgp_outbuf           4104 8192   P         -      -         94  770048
ldp_buffer           4108 8192   P         -      -         7   57344
RPD SNMP             8268 12288  P         -      -         1   12288
LDP config           various          1     896         1     896
-----
349037508      543172620
----- Malloc Usage Report -----
Name           Allocs      Bytes MaxAllocs  MaxBytes  FuncCalls
MGMT.local          1         12         1      12         1
RSVP                 -          -         1     2048    156084
BGP_Group_Tweak-RTClie  2         24         2      24         2
[...Output truncated...]
LDP                 2         24         2      24         2
KRT Request         -          -         1      16    446888
BGP_Group_Packet-Design  2         24         2      24         38
[...Output truncated...]
MPLS                22272  1221656    22274  1221784    228522
BGP.0.0.0.0+179     186419  2237028   192292  2307504  282141191
IS-IS I/O./var/run/ppmd  1     66536      43   103916  695536231

```



```

IS-IS                2407      361372      5887      446076      889294754
BGP RT Background    3        66556        3        66556          3
SNMP Subagent./var/run/ -         24          1        9144        3677022
KRT                  2       205616        3       207900          10
ASPaths             13901     1581544     18023     2067605     293868769
RT                   27         556         28         580         2815
Scheduler            194        2604        199        2684        41382
--Anonymous--        4294944918 4293764616 4294967294 4294967292 45560848
--System--           38565     35474324     38684     35487048     235115763
-----
                                40015436                                41923181
Dynamically allocated memory: 485789696      Maximum: 541736960
  Program data+BSS memory:    2101248      Maximum: 2101248
    Page data overhead:      3039232      Maximum: 3039232
    Page directory size:      512000      Maximum: 512000
-----
Total bytes in use: 491442176 (70% of available memory)

```

Meaning

The sample output shows summary statistics about the entries in the routing table (**show route summary** command) and the memory usage breakdown (**show task memory detail** command) for the routing process (rpd). The two commands provide a comprehensive picture of the memory utilization of the routing protocol process.

The **show route summary** command shows the number of routes in the various routing tables. In the sample output, the routing tables represented are **inet.0**, **inet.2**, **inet.3**, **iso.0**, and **mpls.0**. Within each routing table, all of the active, hold-down, and hidden destinations and routes are summarized for all the protocols from which routes are learned. Routes are in the **hold-down** state prior to being declared inactive, and **hidden** routes are not used because of routing policy. Routes in the **hold-down** and **hidden** states are still using memory because they appear in the routing table.

In addition, routes are summarized in the following categories: those directly connected to the network (**Direct**), local routes (**Local**), and routes learned from configured routing protocols, such as BGP and IS-IS.

The **show task memory detail** command lists the data structures within the tasks run by the routing protocol process (rpd). Tasks are enabled depending on the router's configuration. For example, **isis_area_addr** is a data structure resulting from the IS-IS configuration. The **AllocBytes** field indicates the highest amount of memory used by the data structure. For example, the **isis_area_addr** data structure has 544 blocks of allocated memory, each block is allocated a value of 16 bytes, resulting in allocated bytes of 8704. The maximum allocated blocks and bytes are high-water marks for a data structure. For more information on displaying task-related information, see ["Display Tasks" on page 2446](#).

The **Total bytes in use** field shows the total amount of memory used by the routing protocol process (rpd).

Display Tasks

Purpose

You can display information about tasks to further your investigation of a memory problem on the router.

Action

To display a list of tasks that are enabled on the router, enter the following Junos OS CLI operational mode commands:

```
user@host> show task
user@host> show task memory
user@host> show task task-name
```

Sample Output

```
user@R1> show task
```

```
Pri Task Name                                Pro  Port So Flags
10 LMP Client                                17 <>
10 IF
15 INET6
15 INET
15 ISO
15 Memory
20 RPD Unix Domain Server./var/run/rpd_serv.local 21 <>
20 RPD Unix Domain Server./var/run/rpd_serv.local 20 <>
20 RPD Unix Domain Server./var/run/rpd_serv.local 19 <>
20 RPD Unix Domain Server./var/run/rpd_server_communication 16 <Accept>
20 RPD Server.0.0.0.0+666                      666 15 <Accept>
20 Aggregate
20 RT
30 ICMP                                1
30 Router-Advertisement
30 ICMPv6                                58 9 <>
39 OSPFv2 I/O./var/run/ppmd_control          12 <>
40 l2vpn global task
40 BGP RT Background                      <LowPrio>
40 BGP.::+179                            179 23 <Accept LowPrio>
40 BGP.0.0.0.0+179                        179 22 <Accept LowPrio>
40 BFD I/O./var/run/bfdd_control            11 <>
40 OSPF                                89
```



```

50 BGP_65001.10.0.0.5+3531          3531 18 <LowPrio>
50 BGP_65002.10.1.12.2+1224        1224 25 <LowPrio>
50 BGP_Group_internal               <LowPrio>
50 BGP_Group_toR2                   <LowPrio>
50 TED
50 ASPaths
51 Resolve inet.0                   <LowPrio>
60 KStat                            13 <>
60 KRT Request                      7 <>
60 KRT Ifstate                      255 6 <>
60 KRT                              255 5 <>
60 Redirect
70 MGMT.local                       24 <>
70 MGMT_Listen./var/run/rpd_mgmt    14 <Accept>
70 SNMP Subagent./var/run/snmpd_stream 10 <>

```

80 IF Delete

user@R1> show task memory

Memory	Size (kB)	Percentage	When
Currently In Use:	3490	1%	now
Maximum Ever Used:	3535	1%	04/02/04 11:54:46
Available:	220623	100%	now

user@R1> show task io

Task Name	Reads	Writes	Rcvd	Sent	Dropped
LMP Client	1	1	0	0	0
IF	0	0	0	0	0
INET6	0	0	0	0	0
INET	0	0	0	0	0
ISO	0	0	0	0	0
Memory	0	0	0	0	0
RPD Unix Domain Server./var/ru	1	0	0	0	0
RPD Unix Domain Server./var/ru	1	0	0	0	0
RPD Unix Domain Server./var/ru	0	0	0	0	0
RPD Unix Domain Server./var/ru	3	0	0	0	0
RPD Server.0.0.0.0+666	0	0	0	0	0
Aggregate	0	0	0	0	0
RT	0	0	0	0	0
ICMP	0	0	0	0	0
Router-Advertisement	0	0	0	0	0
ICMPv6	0	0	0	0	0
OSPFv2 I/O./var/run/ppmd_contr	31167	1	0	0	0
l2vpn global task	0	0	0	0	0
BGP RT Background	0	0	0	0	0
BGP.::+179	0	0	0	0	0

BGP.0.0.0.0+179	8	0	0	0	0
BFD I/O./var/run/bfdd_control	30731	1	0	0	0
OSPF	0	0	0	0	0
BGP_65001.10.0.0.5+3531	20486	0	0	0	0
BGP_65002.10.1.12.2+1224	20489	6	0	0	0
BGP_Group_internal	0	0	0	0	0
BGP_Group_toR2	0	0	0	0	0
TED	0	0	0	0	0
ASPaths	0	0	0	0	0
Resolve inet.0	0	0	0	0	0
KStat	0	0	0	0	0
KRT Request	0	0	57	0	0
KRT Ifstate	18	0	16	0	0
KRT	0	0	2	0	0
Redirect	0	0	0	0	0
MGMT.local	0	0	0	0	0
MGMT_Listen./var/run/rpd_mgmt	23	0	0	0	0
SNMP Subagent./var/run/snmpd_s	23	0	0	0	0
IF Delete	0	0	0	0	0

Meaning

The sample output shows a list of routing, routing protocol, and interface tasks that are currently running on the router (**show task**), a summary of memory utilization (**show task memory**), and the memory utilization of a particular task (**show task io**). Tasks can be baseline tasks performed regardless of the router configuration, and other tasks that depend on the router configuration. For example, the **BGP_Group_internal** task is the result of the configuration of BGP on the router, while the **INET6** task is a base task associated with the routing process (rpd).

Each task in the **show task** command output has a priority and a task name. For example, the current priority is 10 for **LMP Client** and 80 for **IF Delete**. A lower number indicates a higher priority.

Some tasks have flags attached to them. For example, the **BGP.0.0.0.0+179** task has two flags, **Accept** and **LowPrio**. The **Accept** flag indicates that the task is waiting for incoming connections, and the **LowPrio** flag indicates that the task will be dispatched to read its socket after other, higher priority tasks. Two additional flags are **Connect**, which indicates that a task is waiting for a connection to complete, and **Delete**, which indicates that a task has been deleted and is being cleaned up.

The **show task io** command shows the statistics gathered for each IO operation. The counters show the following:

- **Reads**—This counter increments when a datagram arrives on a connected socket of the task and the task's read callback is called.
- **Writes**—This counter increments when a connected socket of a task becomes writable and the task's callback is called.
- **Rcvd**—This counter increments when the task calls the Routing Engine to read a datagram from a socket which may or may not be connected.
- **Sent**—This counter increments when a task attempts to read or write a datagram on an existing or nonexisting socket.
- **Drops**—This counter increments when a task attempts to read or write a datagram through the Routing Engine on a prebuilt socket, but the request fails for any reason.

Run Snmpwalk from an NMS System to a Juniper Router

Purpose

Snmpwalk is an SNMP application that you can use to query a MIB for information about the functioning of a router in your network. Snmpwalk uses **GetNext** requests to retrieve the specified information. Object identifiers (OIDs) are used to query the MIB. If the OID argument is not present, Snmpwalk searches MIB-2.

Action

To run Snmpwalk for a specific OID, from a management station that has access to the router, and using a tool such as Snmpwalk, enter the following command:

```
user-nms# snmpwalk [common arguments] hostname community object-id
```

Sample Output

```
user-nms % snmpwalk -Os -M /volume/~mibs -m all tp1 public .1.3.6.1.2.1.4
ipForwarding.0 = forwarding(1)
ipDefaultTTL.0 = 64
ipInReceives.0 = Counter32: 9262713
ipInHdrErrors.0 = Counter32: 0
ipInAddrErrors.0 = Counter32: 0
ipForwDatagrams.0 = Counter32: 614171
ipInUnknownProtos.0 = Counter32: 0
ipInDiscards.0 = Counter32: 0
ipInDelivers.0 = Counter32: 8648408
ipOutRequests.0 = Counter32: 1226483
ipOutDiscards.0 = Counter32: 0
```



```

ipOutNoRoutes.0 = Counter32: 0
ipReasmTimeout.0 = 60
ipReasmReqds.0 = Counter32: 0
ipReasmOKs.0 = Counter32: 0
ipReasmFails.0 = Counter32: 0
ipFragOKs.0 = Counter32: 0
ipFragFails.0 = Counter32: 0
ipFragCreates.0 = Counter32: 0
ipAdEntAddr.10.0.0.1 = IPAddress: 10.0.0.1
ipAdEntAddr.10.1.12.1 = IPAddress: 10.1.12.1
ipAdEntAddr.10.1.15.1 = IPAddress: 10.1.15.1
ipAdEntAddr.10.168.70.143 = IPAddress: 10.168.70.143
[...Output truncated...]

```

Meaning

The sample output shows that the user is on a network management station (**user-nms %**) that has access to the router, **tp1**. In the command, the following options are used:

- **Os**—Deletes all but the last symbolic part of the OID **sysUpTime.0**. For example, Timeticks: (14096763) 1 day, 15:09:27.63.
- **-M**—Compiles the MIB and gives a path or location to the MIBs.
- **-m**—Uses the files in the directory pointed to by the **-M** option.
- **all**—Uses all the files in the directory pointed to by the **-M** option.

In addition, the command includes the hostname **tp1**, the community string **public**, and the OID **.1.3.6.1.2.1.4**.

The OID in this example is from RFC 2096, *IP Forwarding Table MIB*, which displays multipath IP routes that have the same network number but different network masks.

Before you can retrieve SNMP information from a router, you must have the minimum SNMP configuration for that router. Following is the minimum SNMP configuration required:

```

[edit]
snmp {
    community public {
        authorization read-only;
    }
}

```

With this configuration, the system responds to SNMP **Get**, **GetNext**, and **GetBulk** commands that contain the community string **public**.

For more detailed information on configuring SNMP on a router, see the *Junos Network Management Configuration Guide*.

Configure Trace Operations for SNMP

Purpose

Define tracing for SNMP to access more granular information about the packets sent and received through SNMP.

Action

To configure SNMP tracing operations, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@R1# edit snmp
```

2. Configure trace operations:

```
[edit snmp]
user@R1# set traceoptions flag pdu
```

3. Commit the configuration:

```
user@R1# commit and-quit
commit complete
Exiting configuration mode
```

Sample Output

```
user@R1> show configuration snmp
```

```
view all {
  oid .1 include;
}
view system {
  oid system;
}
community public {
  view all;
  authorization read-only;
```



```

}
community private {
    view system;
    authorization read-write;
}
traceoptions {
    flag pdu;
}

```

Meaning

The sample output shows a configuration for SNMP that includes traceoptions. The **pdu** flag is configured, which results in the generation of SNMP request and response packets. The output for the tracing operation is placed into various log files in the **/var/log** directory.

Protocol-specific tracing operations override any equivalent operations that you specify in the global **traceoptions** statement. If there are no equivalent operations, they supplement the global tracing options. If you do not specify any protocol-specific tracing, the routing protocol inherits all the global tracing operations.

Query a MIB With SNMPGet

Purpose

Send an SNMP request to check that the SNMP configuration is correct.

Action

To query a MIB with **SNMPGet**, enter the following command:

```
user@nms % snmpget hostname community oid
```

Sample Output

```

user-nms % snmpget tp1 public .1.3.6.1.2.1.1.0
system.sysDescr.0 = m7i internet router, kernel 6.0R1.5

user-nms % snmpget tp1 public sysDescr.0
system.sysDescr.0 = m7i internet router, kernel 6.0R1.5

```


Meaning

The sample output shows a query from a network management station (**nms**) for the description of the system running on the router **tp1**. The OID is entered in numerical form in the command line, and a description (**sysDescr.0**) is obtained in the output. You can also use **sysDescr.0** in the command line to obtain the same output.

Check CPU Utilization

Purpose

The enterprise-specific chassis MIB provides information about the router and its components. Within the chassis MIB, the **jnxMIBs** branch contains one main subbranch, **jnxBoxAnatomy**, which in turn contains a section, **jnxOperatingTable**. Within **jnxOperatingTable**, and under the **jnxOperatingEntry**, you can use the **jnxOperatingCPU** object to monitor the CPU on your router. (See [Figure 179 on page 2453](#).)

Figure 179: Chassis MIB Tree

```

+--juniperMIB(2636)¶
|  +--jnxProducts(1)¶
|  +--jnxServices(2)¶
|  +--jnxMibs(3)¶
|  |  +--jnxBoxAnatomy(1)¶
|  |  |  +--jnxContainersTable(6) ¶
|  |  |  +--jnxContentsTable(8)¶
|  |  |  +--jnxLEDTTable(10)¶
|  |  |  +--jnxFilledTable(12)¶
|  |  |  +--jnxOperatingTable(13)¶
|  |  |  |  +--jnxOperatingEntry(1)¶
|  |  |  |  |  +-- jnxOperatingContentsIndex(1)¶
|  |  |  |  |  +-- jnxOperatingL1Index(2)¶
|  |  |  |  |  +-- jnxOperatingL2Index(3)¶
|  |  |  |  |  +-- jnxOperatingL3Index(4)¶
|  |  |  |  |  +-- jnxOperatingDescr(5)¶
|  |  |  |  |  +-- jnxOperatingState(6)¶
|  |  |  |  |  +-- jnxOperatingTemp(7)¶
|  |  |  |  |  +-- jnxOperatingCPU(8)¶
|  |  |  |  |  +-- jnxOperatingISR(9)¶
|  |  |  |  |  +-- jnxOperatingDRAMSize(10)¶
|  |  |  |  |  +-- jnxOperatingBuffer(11)¶
|  |  |  |  |  +-- jnxOperatingHeap(12)¶
|  |  |  |  |  +-- jnxOperatingUpTime(13)¶
|  |  |  |  |  +-- jnxOperatingLastRestart(14)¶
|  |  |  |  |  +-- jnxOperatingMemory(15)¶
|  |  |  |  |  +-- jnxOperatingStateOrdered(16)¶
|  |  |  +--jnxRedundancyTable(14)¶
|  |  +--jnxFruTable(15)¶

```

After each object description is a value in parenthesis, such as (1). This value can be used to enter an OID for the specific object. For example, to gather information on the CPU, you can type the object description (**jnxOperatingCPU**) or the OID (**.1.3.6.1.4.1.2636.3.1.13.1.8**).

Action

To check CPU utilization using the Juniper Networks enterprise chassis MIB, from a management station that has access to the router, and using a tool such as Snmpwalk, enter the following command:

```
user-bsd# snmpwalk [common arguments] hostname community object-id
```

Sample Output

```
user-nms % snmpwalk -Os -M /volume/~mibs -m all tp1 public jnxOperatingCPU
jnxOperatingCPU.1.1.1.0 = Gauge32: 0
jnxOperatingCPU.1.1.2.0 = Gauge32: 0
jnxOperatingCPU.1.1.3.0 = Gauge32: 0
jnxOperatingCPU.2.1.0.0 = Gauge32: 0
jnxOperatingCPU.4.1.1.0 = Gauge32: 0
jnxOperatingCPU.4.1.2.0 = Gauge32: 0
jnxOperatingCPU.4.1.3.0 = Gauge32: 0
jnxOperatingCPU.4.1.4.0 = Gauge32: 0
jnxOperatingCPU.6.1.1.0 = Gauge32: 224
jnxOperatingCPU.6.1.2.0 = Gauge32: 224
jnxOperatingCPU.7.1.0.0 = Gauge32: 2
jnxOperatingCPU.7.2.0.0 = Gauge32: 2
jnxOperatingCPU.8.1.1.0 = Gauge32: 0
jnxOperatingCPU.8.2.3.0 = Gauge32: 0
jnxOperatingCPU.8.2.4.0 = Gauge32: 0
jnxOperatingCPU.9.1.0.0 = Gauge32: 6
jnxOperatingCPU.9.1.1.0 = Gauge32: 0

user-nms % snmpwalk -Os -M /volume/~mibs -m all tp1 public jnxOperatingDescr
jnxOperatingDescr.1.1.1.0 = midplane
jnxOperatingDescr.1.1.2.0 = midplane
jnxOperatingDescr.1.1.3.0 = midplane
jnxOperatingDescr.2.1.0.0 = Power Supply A
jnxOperatingDescr.4.1.1.0 = Left Tray front fan
jnxOperatingDescr.4.1.2.0 = Left Tray second fan
jnxOperatingDescr.4.1.3.0 = Left Tray third fan
jnxOperatingDescr.4.1.4.0 = Left Tray fourth fan
jnxOperatingDescr.6.1.1.0 = CFEB Internet Processor Ilv1
jnxOperatingDescr.6.1.2.0 = CFEB Internet Processor Ilv1
jnxOperatingDescr.7.1.0.0 = FPC @ 0/*/*
jnxOperatingDescr.7.2.0.0 = FPC @ 1/*/*
jnxOperatingDescr.8.1.1.0 = PIC: 4x OC-3 SONET, MM @ 0/0/*
```



```
jnxOperatingDescr.8.2.3.0 = PIC: 1x Tunnel @ 1/2/*
jnxOperatingDescr.8.2.4.0 = PIC: 1x G/E, 1000 BASE-SX @ 1/3/*
jnxOperatingDescr.9.1.0.0 = Routing Engine
jnxOperatingDescr.9.1.1.0 = Routing Engine PCMCIA Card
```

Meaning

The sample output shows the percentage CPU utilization on router, **tp1**. The Routing Engine (**9.1.0.0**) has 6 percent CPU utilization, the two CFEB Internet Processors **Ilv1** (**6.1.1.0** and **6.1.2.0**) have 22 percent each, and the FPCs (**7.1.0.0** and **7.2.0.0**) have 2 percent each. Components with a value of zero indicate that the information is either unavailable or inapplicable.

The output for the **jnxOperatingDesc** object provides a description of the separate instances in the **jnxOperatingCPU** object. For example, **9.1.0.0** represents the Routing Engine.

Check CPU Utilization per Process

Purpose

The standard system application MIB (RFC 2287, *Definitions of System-Level Managed Objects for Applications*), describes a set of managed objects that are restricted to information that can be determined from the system itself. The object **sysAppElmtRunCPU** provides information about applications and associated elements that have run or are currently running on the host system. (See [Figure 180 on page 2456](#).)


```

sysAppElmtRunCPU.0.0.116 = Timeticks: (25) 0:00:00.25
sysAppElmtRunCPU.0.0.2023 = Timeticks: (0) 0:00:00.00
sysAppElmtRunCPU.0.0.2131 = Timeticks: (1103) 0:00:11.03
sysAppElmtRunCPU.0.0.2160 = Timeticks: (1599) 0:00:15.99
sysAppElmtRunCPU.0.0.2161 = Timeticks: (4) 0:00:00.04
sysAppElmtRunCPU.0.0.2174 = Timeticks: (1168) 0:00:11.68
sysAppElmtRunCPU.0.0.2324 = Timeticks: (1738) 0:00:17.38
sysAppElmtRunCPU.0.0.16781 = Timeticks: (0) 0:00:00.00
sysAppElmtRunCPU.0.0.18311 = Timeticks: (0) 0:00:00.00
sysAppElmtRunCPU.0.0.26827 = Timeticks: (2) 0:00:00.02
sysAppElmtRunCPU.3.1.1 = Timeticks: (483) 0:00:04.83
sysAppElmtRunCPU.3.2.2163 = Timeticks: (33548776) 3 days, 21:11:27.76
sysAppElmtRunCPU.3.3.2185 = Timeticks: (1314) 0:00:13.14
sysAppElmtRunCPU.3.4.2194 = Timeticks: (5282) 0:00:52.82
sysAppElmtRunCPU.3.7.2168 = Timeticks: (20380) 0:03:23.80
sysAppElmtRunCPU.3.9.2169 = Timeticks: (6703) 0:01:07.03
sysAppElmtRunCPU.3.12.2172 = Timeticks: (337) 0:00:03.37
sysAppElmtRunCPU.3.13.2173 = Timeticks: (36) 0:00:00.36
sysAppElmtRunCPU.3.14.2164 = Timeticks: (39783) 0:06:37.83
sysAppElmtRunCPU.3.15.2175 = Timeticks: (4206) 0:00:42.06
sysAppElmtRunCPU.3.16.2165 = Timeticks: (18) 0:00:00.18
sysAppElmtRunCPU.3.17.2176 = Timeticks: (61) 0:00:00.61
sysAppElmtRunCPU.3.19.2177 = Timeticks: (25) 0:00:00.25
sysAppElmtRunCPU.3.20.2178 = Timeticks: (200) 0:00:02.00
sysAppElmtRunCPU.3.21.2179 = Timeticks: (38) 0:00:00.38
sysAppElmtRunCPU.3.23.2188 = Timeticks: (3175) 0:00:31.75
sysAppElmtRunCPU.3.25.2186 = Timeticks: (44774) 0:07:27.74
sysAppElmtRunCPU.3.26.2180 = Timeticks: (17) 0:00:00.17
sysAppElmtRunCPU.3.27.2181 = Timeticks: (48950) 0:08:09.50
sysAppElmtRunCPU.3.30.2187 = Timeticks: (11) 0:00:00.11
sysAppElmtRunCPU.3.31.2184 = Timeticks: (93) 0:00:00.93
sysAppElmtRunCPU.3.34.2171 = Timeticks: (80) 0:00:00.80
sysAppElmtRunCPU.3.35.2047 = Timeticks: (1585) 0:00:15.85
sysAppElmtRunCPU.3.36.2189 = Timeticks: (30) 0:00:00.30
sysAppElmtRunCPU.3.37.2191 = Timeticks: (326) 0:00:03.26
sysAppElmtRunCPU.5.5.7495 = Timeticks: (24721) 0:04:07.21
sysAppElmtRunCPU.5.6.2167 = Timeticks: (936) 0:00:09.36
sysAppElmtRunCPU.5.6.26829 = Timeticks: (1) 0:00:00.01
sysAppElmtRunCPU.5.8.26828 = Timeticks: (25) 0:00:00.25
sysAppElmtRunCPU.5.28.2182 = Timeticks: (29234) 0:04:52.34
sysAppElmtRunCPU.5.29.2183 = Timeticks: (21) 0:00:00.21

user-nms % snmpwalk -Os -M /~/mibs -m all tp1 public sysAppElmtRunName
sysAppElmtRunName.0.0.0 = (swapper)

```



```

sysAppElmtRunName.0.0.2 = (pagedaemon)
sysAppElmtRunName.0.0.3 = (vmdaemon)
sysAppElmtRunName.0.0.4 = (bufdaemon)
sysAppElmtRunName.0.0.5 = (syncer)
sysAppElmtRunName.0.0.6 = (netdaemon)
sysAppElmtRunName.0.0.7 = (if_pfe)
sysAppElmtRunName.0.0.8 = (if_pfe_listen)
sysAppElmtRunName.0.0.9 = (cb_poll)
sysAppElmtRunName.0.0.10 = (vmuncachedaemon)
sysAppElmtRunName.0.0.11 = (scs_housekeeping)
sysAppElmtRunName.0.0.12 = (if_pic_listen)
sysAppElmtRunName.0.0.116 = mfs
sysAppElmtRunName.0.0.2023 = pccardd
sysAppElmtRunName.0.0.2131 = cron
sysAppElmtRunName.0.0.2160 = /sbin/watchdog
sysAppElmtRunName.0.0.2161 = /usr/sbin/tnetd
sysAppElmtRunName.0.0.2174 = /usr/sbin/tnp.sntpd
sysAppElmtRunName.0.0.2324 = (peer proxy)
sysAppElmtRunName.0.0.16781 = /usr/libexec/getty
sysAppElmtRunName.0.0.18311 = /usr/sbin/xntpd
sysAppElmtRunName.0.0.26827 = telnetd
sysAppElmtRunName.3.1.1 = /sbin/preinit
sysAppElmtRunName.3.2.2163 = /usr/sbin/chassisd
sysAppElmtRunName.3.3.2185 = /usr/sbin/dfwd
sysAppElmtRunName.3.4.2194 = /sbin/dcd
sysAppElmtRunName.3.7.2168 = /usr/sbin/snmpd
sysAppElmtRunName.3.9.2169 = /usr/sbin/mib2d
sysAppElmtRunName.3.12.2172 = /usr/sbin/apspd
sysAppElmtRunName.3.13.2173 = /usr/sbin/vrrpd
sysAppElmtRunName.3.14.2164 = /usr/sbin/alarmd
sysAppElmtRunName.3.15.2175 = /usr/sbin/pfed
sysAppElmtRunName.3.16.2165 = /usr/sbin/craftd
sysAppElmtRunName.3.17.2176 = /usr/sbin/sampled
sysAppElmtRunName.3.19.2177 = /usr/sbin/ilmid
sysAppElmtRunName.3.20.2178 = /usr/sbin/rmopd
sysAppElmtRunName.3.21.2179 = /usr/sbin/cosd
sysAppElmtRunName.3.23.2188 = /usr/sbin/fsad
sysAppElmtRunName.3.25.2186 = /usr/sbin/irsd
sysAppElmtRunName.3.26.2180 = /usr/sbin/nasd
sysAppElmtRunName.3.27.2181 = /usr/sbin/fud
sysAppElmtRunName.3.30.2187 = /usr/sbin/rtsdpd
sysAppElmtRunName.3.31.2184 = /usr/sbin/smartd
sysAppElmtRunName.3.34.2171 = /usr/sbin/inetd
sysAppElmtRunName.3.35.2047 = syslogd

```



```

sysAppElmtRunName.3.36.2189 = /usr/sbin/spd
sysAppElmtRunName.3.37.2191 = /usr/sbin/eccd
sysAppElmtRunName.5.5.7495 = /usr/sbin/rpd
sysAppElmtRunName.5.6.2167 = /usr/sbin/mgd
sysAppElmtRunName.5.6.26829 = mgd: (mgd) (user)/dev/ttyp0
sysAppElmtRunName.5.8.26828 = -cli
sysAppElmtRunName.5.28.2182 = /usr/sbin/ppmd
sysAppElmtRunName.5.29.2183 = /usr/sbin/lmpd

```

Meaning

The sample output shows the number of centi-seconds of total system CPU resources consumed by a particular process. For example, the chassis process (**chassisd**, **3.2.2163**) has consumed 3 days, or 33,548,776 centi-seconds of total system CPU resources.

The **sysAppElmtRunName** object retrieves the name of the OID. For example, **sysAppElmtRunCPU.3.2.2163** represents the chassis process.

Retrieve Version Information about Router Software Components

Purpose

RFC 2790, *Host Resources MIB*, describes a set of managed objects that are useful for managing host systems, including routers.

Action

To retrieve version information about software components on the router, from a management station that has access to the router and using a tool, such as Snmpwalk, enter the following command:

```
user-bsd# snmpwalk [common arguments] hostname community object-id
```

Sample Output

```

user-nms % snmpwalk -Os -M /volume/~/.mibs -m all tp1 public.1.3.6.1.2.1.25.6.3hrSWInstalledIndex.2
= 2
hrSWInstalledIndex.3 = 3
hrSWInstalledIndex.4 = 4
hrSWInstalledIndex.5 = 5
hrSWInstalledIndex.6 = 6

```



```

hrSWInstalledIndex.9 = 9
hrSWInstalledName.2 = "JUNOS Base OS Software Suite [6.0R1.5]"
hrSWInstalledName.3 = "JUNOS Kernel Software Suite [6.0R1.5]"
hrSWInstalledName.4 = "JUNOS Packet Forwarding Engine Support (M7i/M10i) [6.0R1.5]"
hrSWInstalledName.5 = "JUNOS Routing Software Suite [6.0R1.5]"
hrSWInstalledName.6 = "JUNOS Online Documentation [6.0R1.5]"
hrSWInstalledName.9 = "JUNOS Support Tools Package [6.0-20031122-unocM2]"
hrSWInstalledID.2 = OID: zeroDotZero
hrSWInstalledID.3 = OID: zeroDotZero
hrSWInstalledID.4 = OID: zeroDotZero
hrSWInstalledID.5 = OID: zeroDotZero
hrSWInstalledID.6 = OID: zeroDotZero
hrSWInstalledID.9 = OID: zeroDotZero
hrSWInstalledType.2 = operatingSystem(2)
hrSWInstalledType.3 = operatingSystem(2)
hrSWInstalledType.4 = operatingSystem(2)
hrSWInstalledType.5 = operatingSystem(2)
hrSWInstalledType.6 = application(4)
hrSWInstalledType.9 = operatingSystem(2)
hrSWInstalledDate.2 = 2003-8-10,20:34:45.0,-7:0
hrSWInstalledDate.3 = 2003-8-10,20:35:21.0,-7:0
hrSWInstalledDate.4 = 2003-8-10,20:36:30.0,-7:0
hrSWInstalledDate.5 = 2003-8-10,20:36:47.0,-7:0
hrSWInstalledDate.6 = 2003-8-10,20:36:51.0,-7:0
hrSWInstalledDate.9 = 2003-11-22,4:8:47.0,-8:0a1

```

Meaning

The sample output shows the version information for various software components on the router.

Checklist for Displaying Basic Chassis Information

Purpose

[Table 70 on page 2461](#) provides links and commands for displaying basic chassis information, including a list of all Flexible PIC Concentrators (FPCs) and Physical Interface Cards (PICs) installed in the router chassis, the hardware version level, and the serial number.

Action

Table 70: Checklist for Displaying Basic Chassis Information

Task	Command or Action
“Display Basic Chassis Information” on page 2461	<code>show chassis hardware <detail></code>

Display Basic Chassis Information

Purpose

Before you return a router component to Juniper Networks, you must contact the Juniper Networks Technical Assistance Center (JTAC) with the serial number of the failed component and failure information. JTAC will then grant a Return Materials Authorization (RMA).

Action

To display a list of the serial numbers of components installed in the router chassis, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show chassis hardware <detail>
```

Sample Output

```
user@host> show chassis hardware
```

```
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis                               25708         M20
Backplane     REV 03   710-002334  BB9738
Power Supply A REV 06   740-001465  005234         AC
Power Supply B REV 06   740-001465  005237         AC
Display       REV 04   710-001519  BA4681
Routing Engine 0 REV 06   740-003239  1000224893     RE-2.0
Routing Engine 1 REV 06   740-003239  9000022146     RE-2.0
SSB slot 0     REV 02   710-001951  AZ8112         Internet Processor II
SSB slot 1     N/A      N/A         N/A            backup
FPC 0          REV 03   710-003308  BD8455         E-FPC
  PIC 0        REV 08   750-002303  AZ5310         4x F/E, 100 BASE-TX
  PIC 1        REV 07   750-004745  BC9368         2x CT3-NxDS0
FPC 1          REV 03   710-003308  BB9032         E-FPC
  PIC 0        REV 03   750-002914  BC0131         2x OC-3 ATM, MM
```



```
user@host> show chassis hardware
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			00159	M40
Backplane	REV 08	710-000073	AA2125	
Power Supply B	Rev A1	740-000235	000289	DC
Maxicab	REV 08	710-000229	CA4516	
Minicab	REV 04	710-001739	CA4610	
Display	REV 07	710-000150	AA5145	
Routing Engine	REV 07	740-005022	P10865702236	RE-3.0
SCB	REV 03	710-007684	CA3900	Internet Processor II
FPC 1	REV 01	710-001292	AL7435	
PIC 0	REV 03	750-000617	AA3530	1x OC-48 SONET, SMIR
FPC 2	REV 09	710-000175	AA4740	
PIC 0	REV 03	750-000617	AA4557	1x OC-48 SONET, SMIR
FPC 3	REV 01	710-001292	AB4775	
PIC 0	REV 03	750-000612	AA1771	2x OC-3 ATM, MM
PIC 1	REV 03	750-002977	AV3457	2x OC-3 ATM, MM
FPC 5	REV 01	710-001292	AC5118	
PIC 1	REV 03	750-003628	AS8882	1x G/E, 1000 BASE-LH

```
user@host> show chassis hardware detail
```

```
Hardware inventory:
```

Item	Version	Part number	Serial number	Description
Chassis			25708	M20
Backplane	REV 03	710-002334	BB9738	
Power Supply A	REV 06	740-001465	005234	AC
Power Supply B	REV 06	740-001465	005237	AC
Display	REV 04	710-001519	BA4681	
Routing Engine 0	REV 06	740-003239	1000224893	RE-2.0
Routing Engine 0			58000007348d9a01	RE-2.0
Routing Engine 1	REV 06	740-003239	9000022146	RE-2.0
Routing Engine 1			d800000734745701	RE-2.0
SSB slot 0	REV 02	710-001951	AZ8112	Internet Processor II
SSRAM bank 0	REV 02	710-001385	242525	2 Mbytes
SSRAM bank 1	REV 02	710-001385	242741	2 Mbytes
SSRAM bank 2	REV 02	710-001385	242886	2 Mbytes
SSRAM bank 3	REV 02	710-001385	242482	2 Mbytes
SSB slot 1	N/A	N/A	N/A	backup
FPC 0	REV 03	710-003308	BD8455	E-FPC
SSRAM	REV 02	710-001385	241669	2 Mbytes
SDRAM bank 0	REV 01	710-000099	0003409	64 Mbytes

SDRAM bank 1	REV 01	710-000099	0003408	64 Mbytes
PIC 0	REV 08	750-002303	AZ5310	4x F/E, 100 BASE-TX
PIC 1	REV 07	750-004745	BC9368	2x CT3-NxDS0
FPC 1	REV 03	710-003308	BB9032	E-FPC
SSRAM	REV 01	710-001385	V00818	2 Mbytes
SDRAM bank 0	REV 01	710-000099	0003803	64 Mbytes
SDRAM bank 1	REV 01	710-000099	0003847	64 Mbytes
PIC 0	REV 03	750-002914	BC0131	2x OC-3 ATM, MM

Meaning

The sample output is for an M20 and an M40 router. It shows a list of all FPCs and PICs installed in the router chassis, including the hardware version level and serial number.

The **detail** option displays detailed information about hardware, including memory, hardware version level, serial number, and additional information about memory.

If the Routing Engine is identified by a 10- and 16-digit serial number, both numbers are displayed in the output for the **detail** option, and are especially important when processing an RMA for such a Routing Engine. In addition, when you request an RMA for the M40 router, include the **maxicab** serial number.

[Table 71 on page 2463](#) provides a description of all the output fields for the **show chassis hardware** command.

Table 71: Output fields for the show chassis hardware command

Output field	Description
Item	<p>(For M-series routers) Chassis component. Information is displayed about the backplane; power supplies; Routing Engine; maxicab (the connection between the Routing Engine and the backplane, for the M40 router only); System Control Board (SCB), System and Switch Board (SSB), Switching and Forwarding Module (SFM), or Forwarding Engine Board (FEB); Miscellaneous Control Subsystem (MCS) and PFE clock generator (PCG) (for the M160 router only); and each FPC and PIC.</p> <p>(For T-series platforms) Chassis component. Information is displayed about the backplane, power supplies, midplane, Control Board (CB), Connector Interface Panel (CIP), FPC, Front Panel Module (FPM) (craft interface), Power Entry Module (PEM), PIC, SONET Clock Generator (SCG), Small Form-factor Pluggable (SFP) modules, Switch Interface Board (SIB), and Switch Processor Mezzanine Board (SPMB).</p>
Version	Revision level of the chassis component.
Part number	Part number of the chassis component.

Table 71: Output fields for the show chassis hardware command (continued)

Output field	Description
Serial number	Serial number of the chassis component. For all RMAs, the chassis serial number must be provided to JTAC. If the RMA is for the chassis itself, you must obtain the backplane or midplane serial number as well.
Description	Brief description of the hardware item.

NOTE: When you request an RMA, you must also include output from the **show chassis environment** command, the **show version** command, and the troubleshooting output used to identify the failure.

Maintain a Single Configuration File for Both Routing Engines

Purpose

For routers that support multiple Routing Engines, you can specify **re0** and **re1** as group names to ensure that the correct IP addresses are used for each Routing Engine and to maintain a single configuration file for both Routing Engines. It is important that the names of the Routing Engines correspond to a slot position because the names **re0** and **re1** are special group names that you must use for the Routing Engines to recognize which configuration statement to use. Routing Engine 0 must be in slot position 0 and must be named **re0**, and Routing Engine 1 must be in slot position 1 and must be named **re1**.

To maintain a single configuration file for both Routing Engines, follow these steps:

1. [Configure the New Group | 2464](#)
2. [Apply the New Group | 2466](#)

Configure the New Group

Purpose

Each **re0** or **re1** group typically contains, at a minimum, the configuration for the hostname and the management interface (**fxp0**). If each Routing Engine uses a different management interface, the group must also contain the configuration for the backup router and static routes.

Action

To configure the **re0** and **re1** groups, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]  
user@host# edit groups
```

2. Configure the group:

```
[edit groups]  
user@host# set group-name
```

For example:

```
[edit groups]  
user@host# set re0
```

3. To configure the management interface, go to the following hierarchy level:

```
[edit groups]  
user@host# edit groups re0
```

4. Include the following statement:

```
[edit groups re0]  
user@host# set interfaces interface-name unit unit family inet address address
```

For example:

```
[edit groups re0]  
user@host# set interfaces fxp0 unit 0 family inet address 1.1.1.1/24
```

5. To configure the hostname, go to the following hierarchy level:

```
[edit groups re0]  
user@host# edit groups re0 system
```

6. Include the following statement:

```
[edit groups re0 system]  
user@host# set host-name hostname
```

For example:


```
[edit groups re0 system]
user@host# set host-name foo-re0
```

7. Verify the configuration:

```
[edit groups re0]
user@host# show
re0 {
  system {
    host-name foo-re0;
  }
  interfaces {
    fxp0 {
      unit 0 {
        family inet {
          address 1.1.1.1/24;
        }
      }
    }
  }
}
```

8. Commit the configuration:

```
user@host# commit
```

9. Repeat Step 1 through Step 8 for the **re1** group.

Meaning

The sample output in Step 7 shows that the **re0** group contains the minimum configuration for a group, the hostname, and the management interface (**fxp0**). If each Routing Engine uses a different management interface, the group must also contain the configuration for the backup router and static routes.

Apply the New Group

Action

To apply the **re0** group to maintain a single configuration file for both Routing Engines, follow these steps:

1. In configuration mode, go to the top hierarchy level and include the following statement:

```
user@host# [edit]
```



```
user@host# set apply-groups group-name
```

For example:

```
user@host# [edit]
```

```
user@host# set apply-groups re0
```

2. Verify the configuration:

```
user@host# show
```

```
groups {
  re0 {
    system {
      host-name foo-re0;
    }
    interfaces {
      fxp0 {
        unit 0 {
          family inet {
            address 1.1.1.1/24;
          }
        }
      }
    }
  }
  re1 {
    system {
      host-name foo-re1;
    }
    interfaces {
      fxp0 {
        unit 0 {
          family inet {
            address 1.1.1.2/24;
          }
        }
      }
    }
  }
}
```



```
apply-groups [ re0 re1 ];
```

3. Commit the configuration:

```
user@host# commit
```

Meaning

The sample output shows that each group, **re0** and **re1**, has its own IP address that is used for each Routing Engine to maintain a single configuration file.

Configure the New Group

Purpose

Each **re0** or **re1** group typically contains, at a minimum, the configuration for the hostname and the management interface (**fxp0**). If each Routing Engine uses a different management interface, the group must also contain the configuration for the backup router and static routes.

Action

To configure the **re0** and **re1** groups, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit groups
```

2. Configure the group:

```
[edit groups]
user@host# set group-name
```

For example:

```
[edit groups]
user@host# set re0
```

3. To configure the management interface, go to the following hierarchy level:

```
[edit groups]
user@host# edit groups re0
```


4. Include the following statement:

```
[edit groups re0]
user@host# set interfaces interface-name unit unit family inet address address
```

For example:

```
[edit groups re0]
user@host# set interfaces fxp0 unit 0 family inet address 1.1.1.1/24
```

5. To configure the hostname, go to the following hierarchy level:

```
[edit groups re0]
user@host# edit groups re0 system
```

6. Include the following statement:

```
[edit groups re0 system]
user@host# set host-name hostname
```

For example:

```
[edit groups re0 system]
user@host# set host-name foo-re0
```

7. Verify the configuration:

```
[edit groups re0]
user@host# show
re0 {
  system {
    host-name foo-re0;
  }
  interfaces {
    fxp0 {
      unit 0 {
        family inet {
          address 1.1.1.1/24;
        }
      }
    }
  }
}
```


8. Commit the configuration:

```
user@host# commit
```

9. Repeat Step 1 through Step 8 for the **re1** group.

Meaning

The sample output in Step 7 shows that the **re0** group contains the minimum configuration for a group, the hostname, and the management interface (**fxp0**). If each Routing Engine uses a different management interface, the group must also contain the configuration for the backup router and static routes.

Apply the New Group

Action

To apply the **re0** group to maintain a single configuration file for both Routing Engines, follow these steps:

1. In configuration mode, go to the top hierarchy level and include the following statement:

```
user@host# [edit]
```

```
user@host# set apply-groups group-name
```

For example:

```
user@host# [edit]
```

```
user@host# set apply-groups re0
```

2. Verify the configuration:

```
user@host# show
```

```
groups {
  re0 {
    system {
      host-name foo-re0;
    }
    interfaces {
      fxp0 {
```



```

        unit 0 {
            family inet {
                address 1.1.1.1/24;
            }
        }
    }
}
re1 {
    system {
        host-name foo-re1;
    }
    interfaces {
        fxp0 {
            unit 0 {
                family inet {
                    address 1.1.1.2/24;
                }
            }
        }
    }
}
}
}
apply-groups [ re0 re1 ];

```

3. Commit the configuration:

```
user@host# commit
```

Meaning

The sample output shows that each group, **re0** and **re1**, has its own IP address that is used for each Routing Engine to maintain a single configuration file.

Viewing Files and Directories on a Device Running Junos OS

IN THIS SECTION

- [Directories on the Device | 2472](#)
- [Listing Files and Directories | 2473](#)
- [Specifying Filenames and URLs | 2476](#)

Junos OS stores information in files on the device, including configuration files, log files, and device software files. This topic shows some examples of operational commands that you can use to view files and directories on a device running Junos OS.

Sections include:

Directories on the Device

The following table lists some standard directories on a device running Junos OS.

Table 72: Directories on the Device

Directory	Description
<code>/config</code>	This directory is located on the device's router's internal flash drive. It contains the active configuration (juniper.conf) and rollback files 1, 2, and 3.
<code>/var/db/config</code>	This directory is located on the router's device's hard drive and contains rollback files 4 through 49.
<code>/var/tmp</code>	This directory is located on the device's hard drive. It holds core files from the various processes on the Routing Engines. Core files are generated when a particular process crashes and are used by Juniper Networks engineers to diagnose the reason for failure.
<code>/var/log</code>	This directory is located on the device's hard drive. It contains files generated by both the device's logging function as well as the traceoptions command.

Table 72: Directories on the Device (*continued*)

Directory	Description
/var/home	This directory is located on the device's hard drive. It contains a subdirectory for each configured user on the device. These individual user directories are the default file location for many Junos OS commands.
/altroot	This directory is located on the device's hard drive and contains a copy of the root file structure from the internal flash drive. This directory is used in certain disaster recovery modes where the internal flash drive is not operational.
/altconfig	This directory is located on the device's hard drive and contains a copy of the /config file structure from the internal flash drive. This directory is also used in certain disaster recovery modes when the internal flash drive is not operational.

Listing Files and Directories

You can view the device's directory structure as well as individual files by issuing the **file** command in operational mode.

1. To get help about the **file** command, type the following:

```

user@host> file ?
Possible completions:
  <[Enter]>      Execute this command
  archive        Archives files from the system
  checksum       Calculate file checksum
  compare        Compare files
  copy           Copy files (local or remote)
  delete         Delete files from the system
  list           List file information
  rename         Rename files
  show           Show file contents
  source-address Local address to use in originating the connection
  |             Pipe through a command
user@host> file

```

Help shows that the **file** command includes several options for manipulating files.

2. Use the **list** option to see the directory structure of the device. For example, to show the files located in your home directory on the device:

```
user@host> file list
.ssh/
common
```

The default directory for the **file list** command is the home directory of the user logged in to the device. In fact, the user's home directory is the default directory for most of Junos OS commands requiring a filename.

3. To view the contents of other file directories, specify the directory location. For example:

```
user@host> file list /config
juniper.conf
juniper.conf.1.gz
juniper.conf.2.gz
juniper.conf.3.gz
```

4. You can also use the device's context-sensitive help system to locate a directory. For example:

```
user@host> file list /?
Possible completions:
  <[Enter]>          Execute this command
  <path>             Path to list
  /COPYRIGHT         Size: 6355, Last changed: Feb 13 2017
  /altconfig/        Last changed: Aug 07 2017
  /altroot/          Last changed: Aug 07 2017
  /bin/              Last changed: Apr 09 22:31:35
  /boot/             Last changed: Apr 09 23:28:39
  /config/           Last changed: Apr 16 22:35:35
  /data/             Last changed: Aug 07 2017
  /dev/              Last changed: Apr 09 22:36:21
  /etc/              Last changed: Apr 11 03:14:22
  /kernel            Size: 27823246, Last changed: Aug 07 2017
  /mfs/              Last changed: Apr 09 22:36:49
  /mnt/              Last changed: Jan 11 2017
  /modules/          Last changed: Apr 09 22:33:54
  /opt/              Last changed: Apr 09 22:31:00
  /packages/         Last changed: Apr 09 22:34:38
  /proc/             Last changed: May 07 20:25:46
  /rdm.taf           Size: 498, Last changed: Apr 09 22:37:31
  /root/             Last changed: Apr 10 02:19:45
  /sbin/             Last changed: Apr 09 22:33:55
```



```

/staging/           Last changed: Apr 09 23:28:41
/tmp/               Last changed: Apr 11 03:14:49
/usr/               Last changed: Apr 09 22:31:34
/var/               Last changed: Apr 09 22:37:30
user@host> file list /var/?
<[Enter]>           Execute this command
  <path>             Path to list
/var/account/       Last changed: Jul 09 2017
/var/at/            Last changed: Jul 09 2017
/var/backups/        Last changed: Jul 09 2017
/var/bin/           Last changed: Jul 09 2017
/var/crash/         Last changed: Apr 09 22:31:08
/var/cron/          Last changed: Jul 09 2017
/var/db/            Last changed: May 07 20:28:40
/var/empty/         Last changed: Jul 09 2017
/var/etc/           Last changed: Apr 16 22:35:36
/var/heimdal/       Last changed: Jul 10 2017
/var/home/          Last changed: Apr 09 22:59:18
/var/jail/          Last changed: Oct 31 2017
/var/log/           Last changed: Apr 17 02:00:10
/var/mail/          Last changed: Jul 09 2017
/var/messages/      Last changed: Jul 09 2017
/var/named/         Last changed: Jul 10 2017
/var/packages/      Last changed: Jan 18 02:38:59
/var/pdb/           Last changed: Oct 31 2017
/var/preserve/      Last changed: Jul 09 2017
/var/run/           Last changed: Apr 17 02:00:01
/var/rundb/         Last changed: Apr 17 00:46:00
/var/rwho/          Last changed: Jul 09 2017
/var/sdb/           Last changed: Apr 09 22:37:31
/var/spool/         Last changed: Jul 09 2017
/var/sw/            Last changed: Jul 09 2017
/var/tmp/           Last changed: Apr 09 23:28:41
/var/transfer/      Last changed: Jul 09 2017
/var/yp/            Last changed: Jul 09 2017
user@host> file list /var/

```

5. You can also display the contents of a file. For example:

```

user@host>file show /var/log/inventory
Jul  9 23:17:46 CHASSISD release 8.4I0 built by builder on 2017-06-12 07:58:27
UTC
Jul  9 23:18:05 CHASSISD release 8.4I0 built by builder on 2017-06-12 07:58:27

```



```

UTC
Jul  9 23:18:06 Routing Engine 0 - part number 740-003239, serial number
9000016755
Jul  9 23:18:15 Routing Engine 1 - part number 740-003239, serial number
9001018324
Jul  9 23:19:03 SSB 0 - part number 710-001951, serial number AZ8025
Jul  9 23:19:03 SSRAM bank 0 - part number 710-001385, serial number 243071
Jul  9 23:19:03 SSRAM bank 1 - part number 710-001385, serial number 410608
...

```

Specifying Filenames and URLs

In some CLI commands and configuration statements—including **file copy**, **file archive**, **load**, **save**, **set system login user *username* authentication load-key-file**, and **request system software add**—you can include a filename. On a routing matrix, you can include chassis information as part of the filename (for example, **lcc0**, **lcc0-re0**, or **lcc0-re1**).

You can specify a filename or URL in one of the following ways:

- **filename**—File in the user's current directory on the local flash drive. You can use wildcards to specify multiple source files or a single destination file. Wildcards are not supported in Hypertext Transfer Protocol (HTTP) or FTP.

NOTE: Wildcards are supported only by the **file (compare | copy | delete | list | rename | show)** commands. When you issue the **file show** command with a wildcard, it must resolve to one filename.

- **path/filename**—File on the local flash disk.
- **/var/filename** or **/var/path/filename**—File on the local hard disk. You can also specify a file on a local Routing Engine for a specific T640 router on a routing matrix:

```
user@host> file delete lcc0-re0:/var/tmp/junk
```

- **a:filename** or **a:path/filename**—File on the local drive. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.
- **hostname:/path/filename**, **hostname:filename**, **hostname:path/filename**, or **scp://hostname/path/filename**—File on an **scp/ssh** client. This form is not available in the worldwide version of Junos OS. The default path is the user's home directory on the remote system. You can also specify **hostname** as **username@hostname**.

- **ftp://hostname/path/filename**—File on an FTP server. You can also specify **hostname** as **username@hostname** or **username:password@hostname**. The default path is the user's home directory. To specify an absolute path, the path must start with **%2F**; for example, **ftp://hostname/%2Fpath/filename**. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed:

```
user@host> file copy ftp://username@ftp.hostname.net//filename
```

```
file copy ftp.hostname.net: Not logged in.
```

```
user@host> file copy ftp://username:prompt@ftp.hostname.net//filename
```

```
Password for username@ftp.hostname.net:
```

- **http://hostname/path/filename**—File on an HTTP server. You can also specify **hostname** as **username@hostname** or **username:password@hostname**. If a password is required and you omit it, you are prompted for it.
- **re0:/path/filename** or **re1:/path/filename**—File on a local Routing Engine. You can also specify a file on a local Routing Engine for a specific T640 router on a routing matrix:

```
user@host> show log lcc0-re1:chassisd
```

Display File Contents

Purpose

To display the contents of a file on the local router.

Action

To display the contents of a file on the local router, use the following CLI operational mode command:

```
user@host> file show filename
```

Sample Output

```
user@host> file show /var/log/messages
```



```

Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
Apr 13 21:07:36 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:08:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
...Output truncated...

```

Meaning

The sample output shows the contents of the `/var/log/messages` file.

Rename a File on a Router

Action

To rename a file on the local router, use the following CLI operational mode command:

```
user@host> file rename source destination
```

Sample Output

```
user@host> file list /var/tmp
```

```

dcd.core
rpd.core
snmpd.core

```

```
user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
```

```
user@host> file list /var/tmp
```

```

dcd.core.990413
rpd.core
snmpd.core

```

Meaning

The sample output shows that the **dcd.core** file was renamed to **dcd.core.990413**. The original name of the file is the *source* and the new name for the file is the *destination*.

Delete a File on a Router

Action

To delete a file on the local router, use the following CLI operational mode command:

```
user@host> file delete filename
```

Sample Output

```
user@host> file list /var/tmp
```

```
dcd.core  
rpd.core  
snmpd.core
```

```
user@host> file delete /var/tmp/snmpd.core
```

```
user@host> file list /var/tmp
```

```
dcd.core  
rpd.core
```

Meaning

The sample output shows that the **snmpd.core** file was deleted.

Check the Time on a Router

Purpose

Display the current time on a router and display information about how long the router, router software, and routing protocols have been running.

Action

To check time on a router, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> show system uptime
```


Sample Output

```
user@host> show system uptime
```

```
Current time:      1998-10-13 19:45:47 UTC
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

Meaning

The sample output shows the current system time in UTC, the date and time when the router was last booted and how long it has been running, when the routing protocols were last started and how long they have been running, when a configuration was last committed, and the name of the user who issued the last **commit** command. If a different time zone is configured, the output shows that time zone. For information on configuring the time zone, see the *Junos System Basics Configuration Guide*.

The sample output shows that the current time is 12:45 PM, the router has been operational for 22:54 hours, and two users are logged in to the router. The output also shows that the load average is 0.07 seconds for the last minute, 0.02 seconds for the last 5 minutes, and 0.01 seconds for the last 15 minutes.

Check for Users in Configuration Mode

Purpose

Before you change the configuration or commit a candidate configuration, it is a good idea to check for users in configuration mode.

Action

To display users currently editing the configuration, follow these steps:

1. To enter configuration mode, type the following command:

```
user@host> edit
```

2. Enter the following configuration mode command:

```
[edit]
user@host# status
```

For example:


```
user@host> show system users
```

```
4:58PM PST up 5 days, 9:52, 5 users, load averages: 0.01, 0.01, 0.00
USER      TTY      FROM            LOGIN@   IDLE   WHAT
mwazna    p0       bigpunk.juniper.net 4:58PM   -    -cli (cli)
jgchan    p1       bigpunk.juniper.net 2:25PM   2:32 -csh (csh)
jgchan    p2       bigpunk.juniper.net 2:35PM   2:18 cli
taffy     p3       bigpunk.juniper.net 3:28PM   5    -cli (cli)
tmauro    p4       bigpunk.juniper.net 4:16PM   37   cli
```

Sample Output

Meaning

The sample output lists the users who are currently logged in to the router. Five users are logged in to the router, with one user logged in twice, **jgchan**. Each user is logged in through a different terminal (**TTY**—**p0**, **p1**, **p2**, **p3**, and **p4**) from the system **bigpunk.juniper.net**. A hyphen in the **FROM** field indicates that the user logged in through the console.

Additional information includes the time when the user logged in (**LOGIN**), the amount of time the user is not active on the router (**IDLE**), and the processes that the user is running (**WHAT**). In this example, the users are running the command-line interface (**cli**) and the UNIX-level shell (**csh**).

Check the Commands That Users Are Entering

Purpose

A common set of operations you can check is when users log in to the router and the CLI commands they issue.

To check the commands that users are entering, follow these steps:

1. [Configure the Log File for Tracking CLI Commands | 2482](#)
2. [Display the Configured Log File | 2483](#)

Configure the Log File for Tracking CLI Commands

Action

To configure the log file for tracking CLI commands, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit system syslog
```

2. Configure the log file:

```
[edit system syslog]
user@host# edit file filename
```

For example:

```
[edit system syslog]
user@host# edit file cli-commands
```

3. Configure the interactive-commands facility and severity level:

```
[edit system syslog filename]
user@host# set interactive-commands info
```

4. Verify the configuration:

```
[edit system syslog]
user@host# show
file cli-commands {
    interactive-commands info;
}
```

5. Commit the configuration:

```
user@host# commit
```

Meaning

The configuration example shows that the log file **cli-commands** is configured with the **interactive-commands** facility at the **info** severity level. [Table 73 on page 2483](#) lists and describes the severity levels.

Table 73: Severity Levels

Severity Level	Description
info	Log all top-level CLI commands, including the configure command, and all configuration mode commands.
notice	Log the configuration mode commands rollback and commit .
warning	Log when any software process restarts.

Display the Configured Log File

Purpose

To display the log file in configuration mode, enter the following command:

Action

```
[edit system syslog]
user@host# run show log filename
```

For example:

```
[edit system syslog]
user@host# run show log cli-commands
```

Sample Output

```
[edit system syslog]
user@host# run show log cli-commands
Sep 16 11:24:25 nut mgd[3442]: UI_COMMIT_PROGRESS: commit: signaling 'Syslog
daemon', pid 2457, signal 1, status 0
Sep 16 11:24:25 nut mgd[3442]: UI_COMMIT_PROGRESS: commit: signaling 'SNMP daemon',
pid 2592, signal 31, status 0
Sep 16 11:28:36 nut mgd[3442]: UI_CMDLINE_READ_LINE: User 'user', command 'run
show log cli-commands '
Sep 16 11:30:39 nut mgd[3442]: UI_CMDLINE_READ_LINE: User 'user', command 'run
show log security '
Sep 16 11:31:26 nut mgd[3442]: UI_CMDLINE_READ_LINE: User 'user', command 'run
show log messages '
```



```

Sep 16 11:41:21  nut mgd[3442]: UI_CMDLINE_READ_LINE: User 'user', command 'edit
file cli-commands '
Sep 16 11:41:25  nut mgd[3442]: UI_CMDLINE_READ_LINE: User 'user', command 'show
'
Sep 16 11:44:57  nut mgd[3442]: UI_CMDLINE_READ_LINE: User 'user', command 'set
interactive-commands info '
Sep 16 14:32:15  nut mgd[3442]: UI_CMDLINE_READ_LINE: User 'user', command 'run
show log cli-commands '

```

Meaning

The sample output shows the CLI commands that were entered since the log file was configured.

Configure the Log File for Tracking CLI Commands

Action

To configure the log file for tracking CLI commands, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```

[edit]
user@host# edit system syslog

```

2. Configure the log file:

```

[edit system syslog]
user@host# edit file filename

```

For example:

```

[edit system syslog]
user@host# edit file cli-commands

```

3. Configure the interactive-commands facility and severity level:

```

[edit system syslog filename]
user@host# set interactive-commands info

```

4. Verify the configuration:

```

[edit system syslog]

```



```
user@host# show
file cli-commands {
    interactive-commands info;
}
```

5. Commit the configuration:

```
user@host# commit
```

Meaning

The configuration example shows that the log file **cli-commands** is configured with the **interactive-commands** facility at the **info** severity level. [Table 73 on page 2483](#) lists and describes the severity levels.

Table 74: Severity Levels

Severity Level	Description
info	Log all top-level CLI commands, including the configure command, and all configuration mode commands.
notice	Log the configuration mode commands rollback and commit .
warning	Log when any software process restarts.

Check When the Last Configuration Change Occurred

Purpose

When a problem occurs on a router, it is a good idea to check when the last configuration change was made because it may have had some influence on the problem.

To check when the last configuration change occurred, follow these steps:

- 1. [Configure Configuration Change Tracking | 2485](#)
- 2. [Display the Configured Log File | 2486](#)

Configure Configuration Change Tracking

Action

To configure this type of logging, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]  
user@host# edit system syslog
```

2. Configure the log file:

```
[edit system syslog]  
user@host# edit file filename
```

For example:

```
[edit system syslog]  
user@host# edit file mw-configuration-changes
```

3. Configure the change-log facility and severity level:

```
[edit system syslog filename]  
user@host# set change-log info
```

4. Verify the configuration:

```
[edit system syslog]  
user@host# show  
file mw-configuration-changes {  
    change-log info;  
}
```

5. Commit the configuration:

```
user@host# commit
```

Display the Configured Log File

Purpose

To display the log file in configuration mode.

Action

To display the log file in configuration mode, enter the following command:

```
[edit system syslog]
```



```
user@host# run show log filename
```

For example:

```
[edit system syslog]
user@host# run show log mw-configuration-changes
```

Sample Output

```
[edit system syslog]
user@host# run show log mw-configuration-changes
Sep 17 07:03:22  nut mgd[7793]: UI_CFG_AUDIT_OTHER: User 'root' override:
/config/juniper.conf
Sep 17 07:07:21  nut mgd[2751]: UI_CFG_AUDIT_OTHER: User 'root' set: [interfaces
lo0 unit 0 family inet address 127.0.0.1/32]
Sep 17 07:07:21  nut mgd[2751]: UI_CFG_AUDIT_SET: User 'root' set: [system
domain-name] "englab.company.net" -> "englab.company.net"
Sep 17 07:07:21  nut mgd[2751]: UI_CFG_AUDIT_OTHER: User 'root' set: [system
name-server 172.17.28.101]
Sep 17 07:07:22  nut mgd[2751]: UI_CFG_AUDIT_OTHER: User 'root' set: [system
domain-search] "englab.company.net"
Sep 17 07:07:22  nut mgd[2751]: UI_CFG_AUDIT_OTHER: User 'root' set: [system
domain-search] "company.net"
Sep 17 07:07:22  nut mgd[2751]: UI_CFG_AUDIT_OTHER: User 'root' set: [system
domain-search] "jnpr.net"
```

Meaning

The sample output shows the contents of the log file and that the last configuration change was on September 17 at 07:07:22.

Configure Configuration Change Tracking

Action

To configure this type of logging, follow these steps:

1. In configuration mode, go to the following hierarchy level:


```
[edit]
user@host# edit system syslog
```

2. Configure the log file:

```
[edit system syslog]
user@host# edit file filename
```

For example:

```
[edit system syslog]
user@host# edit file mw-configuration-changes
```

3. Configure the change-log facility and severity level:

```
[edit system syslog filename]
user@host# set change-log info
```

4. Verify the configuration:

```
[edit system syslog]
user@host# show
file mw-configuration-changes {
    change-log info;
}
```

5. Commit the configuration:

```
user@host# commit
```

Display a Log File

Purpose

To look at a log or trace file.

Action

To look at a log or trace file, use the following Junos OS CLI operational mode command:

```
user@host> show log filename
```


Sample Output

user@host> **show log messages**

```
Sep 10 07:00:00 host newsyslog[7249]: logfile turned over
Sep 10 07:01:49 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1348
Sep 10 07:04:17 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1349
Sep 10 07:06:45 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1350
Sep 10 07:07:53 host login: 2 LOGIN FAILURES FROM 172.24.16.21
Sep 10 07:07:53 host login: 2 LOGIN FAILURES FROM 172.24.16.21, show configuration
| no-more
Sep 10 07:08:25 host inetd[2785]: /usr/libexec/telnetd[7251]: exit status 0x100
Sep 10 07:09:13 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1351
Sep 10 07:11:41 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1352
Sep 10 07:14:09 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1353
Sep 10 07:16:37 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1354
Sep 10 07:19:05 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor: 10.0.8.1+1355
Sep 10 07:21:33 host rpd[6451]: bgp_listen_accept: Connection attempt from
unconfigured neighbor:
```

Meaning

The sample output shows the **rpd** log messages in the **messages** file for September 10 from 7:00 to 7:21 AM.

NOTE: Local log files are saved in the **/var/log** directory.

Configure IS-IS-Specific Options

Purpose

When unexpected events or problems occur, or if you want to diagnose IS-IS adjacency establishment issues, you can view more detailed information by configuring options specific to IS-IS.

To configure IS-IS options, follow these steps:

1. [Displaying Detailed IS-IS Protocol Information | 2490](#)
2. [Displaying Sent or Received IS-IS Protocol Packets | 2493](#)
3. [Analyzing IS-IS Link-State PDUs in Detail | 2495](#)

Displaying Detailed IS-IS Protocol Information

Action

To trace IS-IS messages in detail, follow these steps:

1. Configure the flag to display detailed IS-IS protocol messages.

```
[edit protocols isis traceoptions]  
user@host# set flag hello detail
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]  
user@host# show  
file isislog size 10k files 10;  
flag hello detail;
```

3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.

```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
```



```

Nov 29 23:17:50 trace_on: Tracing to "/var/log/isislog" started
Nov 29 23:17:50 Sending PTP IIH on so-1/1/1.0
Nov 29 23:17:53 Sending PTP IIH on so-1/1/0.0
Nov 29 23:17:54 Received PTP IIH, source id abc-core-01 on so-1/1/0.0
Nov 29 23:17:54     from interface index 11
Nov 29 23:17:54     max area 0, circuit type 12, packet length 4469
Nov 29 23:17:54     hold time 30, circuit id 6
Nov 29 23:17:54     neighbor state up
Nov 29 23:17:54     speaks IP
Nov 29 23:17:54     area address 99.0008 (1)
Nov 29 23:17:54     IP address 10.10.10.29
Nov 29 23:17:54     4396 bytes of total padding
Nov 29 23:17:54     updating neighbor abc-core-01
Nov 29 23:17:55 Received PTP IIH, source id abc-core-02 on so-1/1/1.0
Nov 29 23:17:55     from interface index 12
Nov 29 23:17:55     max area 0, circuit type 12, packet length 4469
Nov 29 23:17:55     hold time 30, circuit id 6
Nov 29 23:17:55     neighbor state up
Nov 29 23:17:55     speaks IP
Nov 29 23:17:55     area address 99.0000 (1)
Nov 29 23:17:55     IP address 10.10.10.33
Nov 29 23:17:55     4396 bytes of total padding
Nov 29 23:17:55     updating neighbor abc-core-02

```

Meaning

[Table 75 on page 2492](#) lists tracing flags that can be configured specific to IS-IS and presents example output for some of the flags.

Table 75: IS-IS Protocol Tracing Flags

Tracing Flags	Description	Example Output
csn	Complete sequence number PDU (CSNP)	<p>Nov 28 20:02:48 Sending L2 CSN on interface so-1/1/0.0Nov 28 20:02:48 Sending L2 CSN on interface so-1/1/1.0</p> <p>With the detail option.</p> <p>Nov 28 20:06:08 Sending L2 CSN on interface so-1/1/1.0Nov 28 20:06:08 LSP abc-core-01.00-00 lifetime 1146Nov 28 20:06:08 sequence 0x1c4f8 checksum 0xa1e9Nov 28 20:06:08 LSP abc-core-02.00-00 lifetime 411Nov 28 20:06:08 sequence 0x7435 checksum 0x5424Nov 28 20:06:08 LSP abc-brdr-01.00-00 lifetime 465Nov 28 20:06:08 sequence 0xf73 checksum 0xab10Nov 28 20:06:08 LSP abc-edge-01.00-00 lifetime 1089Nov 28 20:06:08 sequence 0x1616 checksum 0xdb29Nov 28 20:06:08 LSP abc-edge-02.00-00 lifetime 1103Nov 28 20:06:08 sequence 0x45cc checksum 0x6883</p>
hello	Hello packet	<p>Nov 28 20:13:50 Sending PTP IIH on so-1/1/1.0Nov 28 20:13:50 Received PTP IIH, source id abc-core-01 on so-1/1/0.0Nov 28 20:13:53 Received PTP IIH, source id abc-core-02 on so-1/1/1.0Nov 28 20:13:57 Sending PTP IIH on so-1/1/0.0Nov 28 20:13:58 Received PTP IIH, source id abc-core-01 on so-1/1/0.0Nov 28 20:13:59 Sending PTP IIH on so-1/1/1.0</p>
lsp	Link-state PDUs (LSPs)	<p>Nov 28 20:15:46 Received L2 LSP abc-edge-01.00-00, interface so-1/1/0.0Nov 28 20:15:46 from abc-core-01Nov 28 20:15:46 sequence 0x1617, checksum 0xd92a, lifetime 1197Nov 28 20:15:46 Updating L2 LSP abc-edge-01.00-00 in TEDNov 28 20:15:47 Received L2 LSP abc-edge-01.00-00, interface so-1/1/1.0Nov 28 20:15:47 from abc-core-02Nov 28 20:15:47 sequence 0x1617, checksum 0xd92a, lifetime 1197</p>
lsp-generation	Link-state PDU generation packets	<p>Nov 28 20:21:24 Regenerating L1 LSP abc-edge-03.00-00, old sequence 0x682Nov 28 20:21:27 Rebuilding L1, fragment abc-edge-03.00-00Nov 28 20:21:27 Rebuilt L1 fragment abc-edge-03.00-00, size 59Nov 28 20:31:52 Regenerating L2 LSP abc-edge-03.00-00, old sequence 0x689Nov 28 20:31:54 Rebuilding L2, fragment abc-edge-03.00-00Nov 28 20:31:54 Rebuilt L2 fragment abc-edge-03.00-00, size 256Nov 28 20:34:05 Regenerating L1 LSP abc-edge-03.00-00, old sequence 0x683Nov 28 20:34:08 Rebuilding L1, fragment abc-edge-03.00-00Nov 28 20:34:08 Rebuilt L1 fragment abc-edge-03.00-00, size 59</p>
packets	All IS-IS protocol packets	Not available.

Table 75: IS-IS Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
psn	Partial sequence number PDU (PSNP) packets	Nov 28 20:40:39 Received L2 PSN, source abc-core-01, interface so-1/1/0.0Nov 28 20:40:39 Received L2 PSN, source abc-core-02, interface so-1/1/1.0Nov 28 20:41:36 Sending L2 PSN on interface so-1/1/1.0Nov 28 20:41:36 Sending L2 PSN on interface so-1/1/0.0Nov 28 20:42:35 Received L2 PSN, source abc-core-02, interface so-1/1/1.0Nov 28 20:42:35 LSP abc-edge-03.00-00 lifetime 1196Nov 28 20:42:35 sequence 0x68c checksum 0x746dNov 28 20:42:35 Received L2 PSN, source abc-core-01, interface so-1/1/0.0Nov 28 20:42:35 LSP abc-edge-03.00-00 lifetime 1196Nov 28 20:42:35 sequence 0x68c checksum 0x746dNov 28 20:42:49 Sending L2 PSN on interface so-1/1/1.0Nov 28 20:42:49 LSP abc-core-01.00-00 lifetime 1197Nov 28 20:42:49 sequence 0x1c4fb checksum 0x9becNov 28 20:42:49 Sending L2 PSN on interface so-1/1/0.0Nov 28 20:42:49 LSP abc-core-01.00-00 lifetime 1197Nov 28 20:42:49 sequence 0x1c4fb checksum 0x9bec
spf	Shortest-path-first (SPF) calculations	Nov 28 20:44:01 Scheduling SPF for L1: ReconfigNov 28 20:44:01 Scheduling multicast SPF for L1: ReconfigNov 28 20:44:01 Scheduling SPF for L2: ReconfigNov 28 20:44:01 Scheduling multicast SPF for L2: ReconfigNov 28 20:44:02 Running L1 SPFNov 28 20:44:02 L1 SPF initialization complete: 0.000099s cumulative timeNov 28 20:44:02 L1 SPF primary processing complete: 0.000303s cumulative timeNov 28 20:44:02 L1 SPF result postprocessing complete: 0.000497s cumulative timeNov 28 20:44:02 L1 SPF RIB postprocessing complete: 0.000626s cumulative timeNov 28 20:44:02 L1 SPF routing table postprocessing complete: 0.000736s cumulative time

SEE ALSO

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

Displaying Sent or Received IS-IS Protocol Packets

To configure the tracing for only sent or received IS-IS protocol packets, follow these steps:

1. Configure the flag to display sent, received, or both sent and received packets.

```
[edit protocols isis traceoptions]
user@host# set flag hello send
```


or

```
[edit protocols isis traceoptions]  
user@host# set flag hello receive
```

or

```
[edit protocols isis traceoptions]  
user@host# set flag hello
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]  
user@host# show  
file isislog size 10k files 10;  
flag hello send;
```

or

```
[edit protocols isis traceoptions]  
user@host# show  
file isislog size 10k files 10;  
flag hello receive;
```

or

```
[edit protocols isis traceoptions]  
user@host# show  
file isislog size 10k files 10;  
flag hello send receive;
```

3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.


```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
Sep 27 18:17:01 ISIS periodic xmit to 01:80:c2:00:00:15 (IFL 2)
Sep 27 18:17:01 ISIS periodic xmit to 01:80:c2:00:00:14 (IFL 2)
Sep 27 18:17:03 ISIS periodic xmit to 01:80:c2:00:00:15 (IFL 2)
Sep 27 18:17:04 ISIS periodic xmit to 01:80:c2:00:00:14 (IFL 2)
Sep 27 18:17:06 ISIS L2 hello from 0000.0000.0008 (IFL 2) absorbed
Sep 27 18:17:06 ISIS periodic xmit to 01:80:c2:00:00:15 (IFL 2)
Sep 27 18:17:06 ISIS L1 hello from 0000.0000.0008 (IFL 2) absorbed
```

SEE ALSO

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

Analyzing IS-IS Link-State PDUs in Detail

To analyze IS-IS link-state PDUs in detail, follow these steps:

1. Configure IS-IS open messages.

```
[edit protocols isis traceoptions]
user@host# set flag lsp detail
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]
user@host# show
file isislog size 5m world-readable;
flag error;
flag lsp detail;
```


3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.

```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
Nov 28 20:17:24 Received L2 LSP abc-core-01.00-00, interface so-1/1/0.0
Nov 28 20:17:24      from abc-core-01
Nov 28 20:17:24      sequence 0x1c4f9, checksum 0x9fea, lifetime 1199
Nov 28 20:17:24      max area 0, length 426
Nov 28 20:17:24      no partition repair, no database overload
Nov 28 20:17:24      IS type 3, metric type 0
Nov 28 20:17:24      area address 99.0908 (1)
Nov 28 20:17:24      speaks CLNP
Nov 28 20:17:24      speaks IP
Nov 28 20:17:24      dyn hostname abc-core-01
Nov 28 20:17:24      IP address 10.10.134.11
Nov 28 20:17:24      IP prefix: 10.10.10.0/30 metric 1 up
Nov 28 20:17:24      IP prefix: 10.10.10.4/30 metric 5 up
Nov 28 20:17:24      IP prefix: 10.10.10.56/30 metric 5 up
Nov 28 20:17:24      IP prefix: 10.10.10.52/30 metric 1 up
Nov 28 20:17:24      IP prefix: 10.10.10.64/30 metric 5 up
Nov 28 20:17:24      IP prefix: 10.10.10.20/30 metric 5 up
Nov 28 20:17:24      IP prefix: 10.10.10.28/30 metric 5 up
Nov 28 20:17:24      IP prefix: 10.10.10.44/30 metric 5 up
Nov 28 20:17:24      IP prefix 10.10.10.0 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 1
Nov 28 20:17:24      IP prefix 10.10.10.4 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.10.10.56 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.10.10.52 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 1
Nov 28 20:17:24      IP prefix 10.10.10.64 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.10.10.20 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.10.10.28 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
```



```

Nov 28 20:17:24      IP prefix 10.10.10.44 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbors:
Nov 28 20:17:24      IS neighbor abc-core-02.00
Nov 28 20:17:24      internal, metrics: default 1
[...Output truncated...]
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbor abc-brdr-01.00
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbor abc-core-02.00, metric: 1
Nov 28 20:17:24      IS neighbor abc-esr-02.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-03.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-01.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-02.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-brdr-01.00, metric: 5
Nov 28 20:17:24      IP prefix: 10.10.134.11/32 metric 0 up
Nov 28 20:17:24      IP prefix: 10.11.0.0/16 metric 5 up
Nov 28 20:17:24      IP prefix: 10.211.0.0/16 metric 0 up
Nov 28 20:17:24      IP prefix 10.10.134.11 255.255.255.255
Nov 28 20:17:24      internal, metrics: default 0
Nov 28 20:17:24      IP prefix 10.11.0.0 255.255.0.0
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.211.0.0 255.255.0.0
Nov 28 20:17:24      internal, metrics: default 0
Nov 28 20:17:24      Updating LSP
Nov 28 20:17:24 Updating L2 LSP abc-core-01.00-00 in TED
Nov 28 20:17:24 Analyzing subtlv's for abc-core-02.00
Nov 28 20:17:24 Analysis complete
Nov 28 20:17:24 Analyzing subtlv's for abc-esr-02.00
Nov 28 20:17:24 Analysis complete
Nov 28 20:17:24 Analyzing subtlv's for abc-edge-03.00
Nov 28 20:17:24 Analysis complete
Nov 28 20:17:24 Analyzing subtlv's for abc-edge-01.00
Nov 28 20:17:24 Analysis complete
Nov 28 20:17:24 Analyzing subtlv's for abc-edge-02.00
Nov 28 20:17:24 Analysis complete
Nov 28 20:17:24 Analyzing subtlv's for abc-brdr-01.00
Nov 28 20:17:24 Analysis complete
Nov 28 20:17:24      Scheduling L2 LSP abc-core-01.00-00 sequence 0x1c4f9 on
interface so-1/1/1.0

```

SEE ALSO

Displaying Detailed IS-IS Protocol Information

Action

To trace IS-IS messages in detail, follow these steps:

1. Configure the flag to display detailed IS-IS protocol messages.

```
[edit protocols isis traceoptions]  
user@host# set flag hello detail
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]  
user@host# show  
file isislog size 10k files 10;  
flag hello detail;
```

3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.

```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
```

```
Nov 29 23:17:50 trace_on: Tracing to "/var/log/isislog" started  
Nov 29 23:17:50 Sending PTP IIH on so-1/1/1.0  
Nov 29 23:17:53 Sending PTP IIH on so-1/1/0.0  
Nov 29 23:17:54 Received PTP IIH, source id abc-core-01 on so-1/1/0.0
```



```

Nov 29 23:17:54      from interface index 11
Nov 29 23:17:54      max area 0, circuit type 12, packet length 4469
Nov 29 23:17:54      hold time 30, circuit id 6
Nov 29 23:17:54      neighbor state up
Nov 29 23:17:54      speaks IP
Nov 29 23:17:54      area address 99.0008 (1)
Nov 29 23:17:54      IP address 10.10.10.29
Nov 29 23:17:54      4396 bytes of total padding
Nov 29 23:17:54      updating neighbor abc-core-01
Nov 29 23:17:55 Received PTP IIH, source id abc-core-02 on so-1/1/1.0
Nov 29 23:17:55      from interface index 12
Nov 29 23:17:55      max area 0, circuit type 12, packet length 4469
Nov 29 23:17:55      hold time 30, circuit id 6
Nov 29 23:17:55      neighbor state up
Nov 29 23:17:55      speaks IP
Nov 29 23:17:55      area address 99.0000 (1)
Nov 29 23:17:55      IP address 10.10.10.33
Nov 29 23:17:55      4396 bytes of total padding
Nov 29 23:17:55      updating neighbor abc-core-02

```

Meaning

[Table 75 on page 2492](#) lists tracing flags that can be configured specific to IS-IS and presents example output for some of the flags.

Table 76: IS-IS Protocol Tracing Flags

Tracing Flags	Description	Example Output
csn	Complete sequence number PDU (CSNP)	<p>Nov 28 20:02:48 Sending L2 CSN on interface so-1/1/0.0Nov 28 20:02:48 Sending L2 CSN on interface so-1/1/1.0</p> <p>With the detail option.</p> <p>Nov 28 20:06:08 Sending L2 CSN on interface so-1/1/1.0Nov 28 20:06:08 LSP abc-core-01.00-00 lifetime 1146Nov 28 20:06:08 sequence 0x1c4f8 checksum 0xa1e9Nov 28 20:06:08 LSP abc-core-02.00-00 lifetime 411Nov 28 20:06:08 sequence 0x7435 checksum 0x5424Nov 28 20:06:08 LSP abc-brdr-01.00-00 lifetime 465Nov 28 20:06:08 sequence 0xf73 checksum 0xab10Nov 28 20:06:08 LSP abc-edge-01.00-00 lifetime 1089Nov 28 20:06:08 sequence 0x1616 checksum 0xdb29Nov 28 20:06:08 LSP abc-edge-02.00-00 lifetime 1103Nov 28 20:06:08 sequence 0x45cc checksum 0x6883</p>

Table 76: IS-IS Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
hello	Hello packet	Nov 28 20:13:50 Sending PTP IIH on so-1/1/1.0Nov 28 20:13:50 Received PTP IIH, source id abc-core-01 on so-1/1/0.0Nov 28 20:13:53 Received PTP IIH, source id abc-core-02 on so-1/1/1.0Nov 28 20:13:57 Sending PTP IIH on so-1/1/0.0Nov 28 20:13:58 Received PTP IIH, source id abc-core-01 on so-1/1/0.0Nov 28 20:13:59 Sending PTP IIH on so-1/1/1.0
lsp	Link-state PDUs (LSPs)	Nov 28 20:15:46 Received L2 LSP abc-edge-01.00-00, interface so-1/1/0.0Nov 28 20:15:46 from abc-core-01Nov 28 20:15:46 sequence 0x1617, checksum 0xd92a, lifetime 1197Nov 28 20:15:46 Updating L2 LSP abc-edge-01.00-00 in TEDNov 28 20:15:47 Received L2 LSP abc-edge-01.00-00, interface so-1/1/1.0Nov 28 20:15:47 from abc-core-02Nov 28 20:15:47 sequence 0x1617, checksum 0xd92a, lifetime 1197
lsp-generation	Link-state PDU generation packets	Nov 28 20:21:24 Regenerating L1 LSP abc-edge-03.00-00, old sequence 0x682Nov 28 20:21:27 Rebuilding L1, fragment abc-edge-03.00-00Nov 28 20:21:27 Rebuilt L1 fragment abc-edge-03.00-00, size 59Nov 28 20:31:52 Regenerating L2 LSP abc-edge-03.00-00, old sequence 0x689Nov 28 20:31:54 Rebuilding L2, fragment abc-edge-03.00-00Nov 28 20:31:54 Rebuilt L2 fragment abc-edge-03.00-00, size 256Nov 28 20:34:05 Regenerating L1 LSP abc-edge-03.00-00, old sequence 0x683Nov 28 20:34:08 Rebuilding L1, fragment abc-edge-03.00-00Nov 28 20:34:08 Rebuilt L1 fragment abc-edge-03.00-00, size 59
packets	All IS-IS protocol packets	Not available.
psn	Partial sequence number PDU (PSNP) packets	Nov 28 20:40:39 Received L2 PSN, source abc-core-01, interface so-1/1/0.0Nov 28 20:40:39 Received L2 PSN, source abc-core-02, interface so-1/1/1.0Nov 28 20:41:36 Sending L2 PSN on interface so-1/1/1.0Nov 28 20:41:36 Sending L2 PSN on interface so-1/1/0.0Nov 28 20:42:35 Received L2 PSN, source abc-core-02, interface so-1/1/1.0Nov 28 20:42:35 LSP abc-edge-03.00-00 lifetime 1196Nov 28 20:42:35 sequence 0x68c checksum 0x746dNov 28 20:42:35 Received L2 PSN, source abc-core-01, interface so-1/1/0.0Nov 28 20:42:35 LSP abc-edge-03.00-00 lifetime 1196Nov 28 20:42:35 sequence 0x68c checksum 0x746dNov 28 20:42:49 Sending L2 PSN on interface so-1/1/1.0Nov 28 20:42:49 LSP abc-core-01.00-00 lifetime 1197Nov 28 20:42:49 sequence 0x1c4fb checksum 0x9becNov 28 20:42:49 Sending L2 PSN on interface so-1/1/0.0Nov 28 20:42:49 LSP abc-core-01.00-00 lifetime 1197Nov 28 20:42:49 sequence 0x1c4fb checksum 0x9bec

Table 76: IS-IS Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
spf	Shortest-path-first (SPF) calculations	<p>Nov 28 20:44:01 Scheduling SPF for L1: ReconfigNov 28 20:44:01 Scheduling multicast SPF for L1: ReconfigNov 28 20:44:01 Scheduling SPF for L2: ReconfigNov 28 20:44:01 Scheduling multicast SPF for L2: ReconfigNov 28 20:44:02 Running L1 SPFNov 28 20:44:02 L1 SPF initialization complete: 0.000099s cumulative timeNov 28 20:44:02 L1 SPF primary processing complete: 0.000303s cumulative timeNov 28 20:44:02 L1 SPF result postprocessing complete: 0.000497s cumulative timeNov 28 20:44:02 L1 SPF RIB postprocessing complete: 0.000626s cumulative timeNov 28 20:44:02 L1 SPF routing table postprocessing complete: 0.000736s cumulative time</p>

RELATED DOCUMENTATION

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

Analyzing IS-IS Link-State PDUs in Detail

To analyze IS-IS link-state PDUs in detail, follow these steps:

1. Configure IS-IS open messages.

```
[edit protocols isis traceoptions]
user@host# set flag lsp detail
```

2. Verify the configuration.

```
user@host# show
```

For example:

```
[edit protocols isis traceoptions]
user@host# show
file isislog size 5m world-readable;
```



```
flag error;
flag lsp detail;
```

3. Commit the configuration.

```
user@host# commit
```

4. View the contents of the file containing the detailed messages.

```
user@host# run show log filename
```

For example:

```
user@host# run show log isislog
Nov 28 20:17:24 Received L2 LSP abc-core-01.00-00, interface so-1/1/0.0
Nov 28 20:17:24     from abc-core-01
Nov 28 20:17:24     sequence 0x1c4f9, checksum 0x9fea, lifetime 1199
Nov 28 20:17:24     max area 0, length 426
Nov 28 20:17:24     no partition repair, no database overload
Nov 28 20:17:24     IS type 3, metric type 0
Nov 28 20:17:24     area address 99.0908 (1)
Nov 28 20:17:24     speaks CLNP
Nov 28 20:17:24     speaks IP
Nov 28 20:17:24     dyn hostname abc-core-01
Nov 28 20:17:24     IP address 10.10.134.11
Nov 28 20:17:24     IP prefix: 10.10.10.0/30 metric 1 up
Nov 28 20:17:24     IP prefix: 10.10.10.4/30 metric 5 up
Nov 28 20:17:24     IP prefix: 10.10.10.56/30 metric 5 up
Nov 28 20:17:24     IP prefix: 10.10.10.52/30 metric 1 up
Nov 28 20:17:24     IP prefix: 10.10.10.64/30 metric 5 up
Nov 28 20:17:24     IP prefix: 10.10.10.20/30 metric 5 up
Nov 28 20:17:24     IP prefix: 10.10.10.28/30 metric 5 up
Nov 28 20:17:24     IP prefix: 10.10.10.44/30 metric 5 up
Nov 28 20:17:24     IP prefix 10.10.10.0 255.255.255.252
Nov 28 20:17:24     internal, metrics: default 1
Nov 28 20:17:24     IP prefix 10.10.10.4 255.255.255.252
Nov 28 20:17:24     internal, metrics: default 5
Nov 28 20:17:24     IP prefix 10.10.10.56 255.255.255.252
Nov 28 20:17:24     internal, metrics: default 5
Nov 28 20:17:24     IP prefix 10.10.10.52 255.255.255.252
Nov 28 20:17:24     internal, metrics: default 1
Nov 28 20:17:24     IP prefix 10.10.10.64 255.255.255.252
```



```

Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.10.10.20 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.10.10.28 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.10.10.44 255.255.255.252
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbors:
Nov 28 20:17:24      IS neighbor abc-core-02.00
Nov 28 20:17:24      internal, metrics: default 1
[...Output truncated...]
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbor abc-brdr-01.00
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IS neighbor abc-core-02.00, metric: 1
Nov 28 20:17:24      IS neighbor abc-esr-02.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-03.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-01.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-edge-02.00, metric: 5
Nov 28 20:17:24      IS neighbor abc-brdr-01.00, metric: 5
Nov 28 20:17:24      IP prefix: 10.10.134.11/32 metric 0 up
Nov 28 20:17:24      IP prefix: 10.11.0.0/16 metric 5 up
Nov 28 20:17:24      IP prefix: 10.211.0.0/16 metric 0 up
Nov 28 20:17:24      IP prefix 10.10.134.11 255.255.255.255
Nov 28 20:17:24      internal, metrics: default 0
Nov 28 20:17:24      IP prefix 10.11.0.0 255.255.0.0
Nov 28 20:17:24      internal, metrics: default 5
Nov 28 20:17:24      IP prefix 10.211.0.0 255.255.0.0
Nov 28 20:17:24      internal, metrics: default 0
Nov 28 20:17:24      Updating LSP
Nov 28 20:17:24      Updating L2 LSP abc-core-01.00-00 in TED
Nov 28 20:17:24      Analyzing subtlv's for abc-core-02.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-esr-02.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-edge-03.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-edge-01.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-edge-02.00
Nov 28 20:17:24      Analysis complete
Nov 28 20:17:24      Analyzing subtlv's for abc-brdr-01.00
Nov 28 20:17:24      Analysis complete

```



```
Nov 28 20:17:24      Scheduling L2 LSP abc-core-01.00-00 sequence 0x1c4f9 on
interface so-1/1/1.0
```

RELATED DOCUMENTATION

Understanding IS-IS Areas to Divide an Autonomous System into Smaller Groups

Example: Configuring a Multi-Level IS-IS Topology to Control Interarea Flooding

Configure OSPF-Specific Options

Purpose

When unexpected events or problems occur, or if you want to diagnose OSPF neighbor establishment issues, you can view more detailed information by configuring options specific to OSPF.

To configure OSPF options, follow these steps:

1. [Diagnose OSPF Session Establishment Problems | 2504](#)
2. [Analyze OSPF Link-State Advertisement Packets in Detail | 2509](#)

Diagnose OSPF Session Establishment Problems

Action

To trace OSPF messages in detail, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols ospf traceoptions
```

2. Configure OSPF hello messages:

```
[edit protocols ospf traceoptions]
user@host# set flag hello detail
```

3. Verify the configuration:

```
user@host# show
```


For example:

```
[edit protocols ospf traceoptions]
user@host# show
file ospf size 5m world-readable;
flag hello detail;
```

4. Commit the configuration:

```
user@host# commit
```

5. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
user@host# run show log ospf
```

```
Dec 2 16:14:24 Version 2, length 44, ID 10.0.0.6, area 1.0.0.0
Dec 2 16:14:24 checksum 0xf01a, authtype 0
Dec 2 16:14:24 mask 0.0.0.0, hello_ivl 10, opts 0x2, prio 128
Dec 2 16:14:24 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
Dec 2 16:14:24 OSPF sent Hello (1) -> 224.0.0.5 (so-1/1/2.0)
Dec 2 16:14:24 Version 2, length 44, ID 10.0.0.6, area 1.0.0.0
Dec 2 16:14:24 checksum 0xf01a, authtype 0
Dec 2 16:14:24 mask 0.0.0.0, hello_ivl 10, opts 0x2, prio 128
Dec 2 16:14:24 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
Dec 2 16:14:26 OSPF rcvd Hello 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0)
Dec 2 16:14:26 Version 2, length 48, ID 10.10.134.12, area 0.0.0.0
Dec 2 16:14:26 checksum 0x99b8, authtype 0Dec 2 16:14:26 mask
255.255.255.252, hello_ivl 10, opts 0x2, prio 1
ec 2 16:14:26 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
Dec 2 16:14:29 OSPF rcvd Hello 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0)
Dec 2 16:14:29 Version 2, length 48, ID 10.108.134.11, area 0.0.0.0
Dec 2 16:14:29 checksum 0x99b9, authtype 0Dec 2 16:14:29 mask
255.255.255.252, hello_ivl 10, opts 0x2, prio 1
Dec 2 16:14:29 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
```

Meaning

[Table 77 on page 2506](#) lists OSPF tracing flags and presents example output for some of the flags.

Table 77: OSPF Protocol Tracing Flags

Tracing Flags	Description	Example Output
database description	All database description packets	Dec 2 15:44:51 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.29 (so-1/1/0.0) state changed from Full to Down Dec 2 15:44:51 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Full to Down Dec 2 15:44:55 RPD_OSPF_NBRUP: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Init to ExStart Dec 2 15:44:55 OSPF sent DbD (2) -> 224.0.0.5 (so-1/1/1.0) Dec 2 15:44:55 Version 2, length 32, ID 10.0.0.6, area 0.0.0.0 Dec 2 15:44:55 checksum 0xf76b, authtype 0 Dec 2 15:44:55 options 0x42, i 1, m 1, ms 1, seq 0xa009eee, mtu 4470 Dec 2 15:44:55 OSPF rcvd DbD 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0) Dec 2 15:44:55 Version 2, length 32, ID 10.10.134.12, area 0.0.0.0 Dec 2 15:44:55 checksum 0x312c, authtype 0 Dec 2 15:44:55 options 0x42, i 1, m 1, ms 1, seq 0x2154, mtu 4470
error	OSPF errored packets	Dec 2 15:49:34 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:49:44 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:49:54 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:50:04 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:50:14 OSPF packet ignored: no matching interface from 172.16.120.29

Table 77: OSPF Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
event	OSPF state transitions	<p>Dec 2 15:52:35 OSPF interface ge-2/2/0.0 state changed from DR to DR</p> <p>Dec 2 15:52:35 OSPF interface ge-3/1/0.0 state changed from DR to DR</p> <p>Dec 2 15:52:35 OSPF interface ge-3/2/0.0 state changed from DR to DR</p> <p>Dec 2 15:52:35 OSPF interface ge-4/2/0.0 state changed from DR to DR</p> <p>Dec 2 15:53:21 OSPF neighbor 10.10.10.29 (so-1/1/0.0) state changed from Full to Down</p> <p>Dec 2 15:53:21 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.29 (so-1/1/0.0) state changed from Full to Down</p> <p>Dec 2 15:53:21 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Full to Down</p> <p>Dec 2 15:53:21 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Full to Down</p> <p>Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Down to Init</p> <p>Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Init to ExStart</p> <p>Dec 2 15:53:25 RPD_OSPF_NBRUP: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Init to ExStart</p> <p>Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from ExStart to Exchange</p> <p>Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Exchange to Full</p> <p>Dec 2 15:53:25 RPD_OSPF_NBRUP: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Exchange to Full</p>
flooding	Link-state flooding packets	<p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 flooding on so-1/1/0.0</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 flooding on so-1/1/1.0</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 on no so-1/1/2.0 retransmit lists, no flood</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 on no so-1/1/3.0 retransmit lists, no flood</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.245.0.1 10.0.0.6 on no so-1/1/2.0 retransmit lists, no flood</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.245.0.1 10.0.0.6 on no so-1/1/3.0 retransmit lists, no flood</p>

Table 77: OSPF Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
hello	Hello packets	Dec 2 15:57:25 OSPF sent Hello (1) -> 224.0.0.5 (ge-3/1/0.0) Dec 2 15:57:25 Version 2, length 44, ID 10.0.0.6, area 2.0.0.0 Dec 2 15:57:25 checksum 0xe43f, authtype 0 Dec 2 15:57:25 mask 255.255.0.0, hello_ivl 10, opts 0x2, prio 128 Dec 2 15:57:25 dead_ivl 40, DR 10.218.0.1, BDR 0.0.0.0 Dec 2 15:57:25 OSPF rcvd Hello 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0) Dec 2 15:57:25 Version 2, length 48, ID 10.10.134.12, area 0.0.0.0 Dec 2 15:57:25 checksum 0x99b8, authtype 0 Dec 2 15:57:25 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 1 Dec 2 15:57:25 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0 Dec 2 15:57:27 OSPF sent Hello (1) -> 224.0.0.5 (ge-3/2/0.0) Dec 2 15:57:27 Version 2, length 44, ID 10.0.0.6, area 2.0.0.0 Dec 2 15:57:27 checksum 0xe4a5, authtype 0 Dec 2 15:57:27 mask 255.255.0.0, hello_ivl 10, opts 0x2, prio 128 Dec 2 15:57:27 dead_ivl 40, DR 10.116.0.1, BDR 0.0.0.0 Dec 2 15:57:28 OSPF rcvd Hello 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0) Dec 2 15:57:28 Version 2, length 48, ID 10.10.134.11, area 0.0.0.0 Dec 2 15:57:28 checksum 0x99b9, authtype 0 Dec 2 15:57:28 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 1 Dec 2 15:57:28 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
lsa-ack	Link-state acknowledgment packets	Dec 2 16:00:11 OSPF rcvd LSAck 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0) Dec 2 16:00:11 Version 2, length 44, ID 10.10.134.11, area 0.0.0.0 Dec 2 16:00:11 checksum 0xcdbf, authtype 0 Dec 2 16:00:11 OSPF rcvd LSAck 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0) Dec 2 16:00:11 Version 2, length 144, ID 10.10.134.12, area 0.0.0.0 Dec 2 16:00:11 checksum 0x73bc, authtype 0 Dec 2 16:00:16 OSPF rcvd LSAck 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0) Dec 2 16:00:16 Version 2, length 44, ID 10.10.134.12, area 0.0.0.0 Dec 2 16:00:16 checksum 0x8180, authtype 0
lsa-request	Link-state request packets	Dec 2 16:01:38 OSPF rcvd LSReq 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0) Dec 2 16:01:38 Version 2, length 108, ID 10.10.134.11, area 0.0.0.0 Dec 2 16:01:38 checksum 0xe86, authtype 0

Table 77: OSPF Protocol Tracing Flags *(continued)*

Tracing Flags	Description	Example Output
lsa-update	Link-state update packets	Dec 2 16:09:12 OSPF built router LSA, area 0.0.0.0 Dec 2 16:09:12 OSPF built router LSA, area 1.0.0.0 Dec 2 16:09:12 OSPF built router LSA, area 2.0.0.0 Dec 2 16:09:13 OSPF sent LSUpdate (4) -> 224.0.0.5 (so-1/1/0.0) Dec 2 16:09:13 Version 2, length 268, ID 10.0.0.6, area 0.0.0.0 Dec 2 16:09:13 checksum 0x8047, authtype 0 Dec 2 16:09:13 adv count 7 Dec 2 16:09:13 OSPF sent LSUpdate (4) -> 224.0.0.5 (so-1/1/1.0) Dec 2 16:09:13 Version 2, length 268, ID 10.0.0.6, area 0.0.0.0 Dec 2 16:09:13 checksum 0x8047, authtype 0 Dec 2 16:09:13 adv count 7
packets	All OSPF packets	Not available.
packet-dump	Dump the contents of selected packet types	Not available.
spf	SPF calculations	Dec 2 16:08:03 OSPF full SPF refresh scheduled Dec 2 16:08:04 OSPF SPF start, area 1.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.0.0.6 distance 0 to SPF list Dec 2 16:08:04 SPF elapsed time 0.000525s Dec 2 16:08:04 Stub elapsed time 0.000263s Dec 2 16:08:04 OSPF SPF start, area 2.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.0.0.6 distance 0 to SPF list Dec 2 16:08:04 SPF elapsed time 0.000253s Dec 2 16:08:04 Stub elapsed time 0.000249s Dec 2 16:08:04 OSPF SPF start, area 0.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.0.0.6 distance 0 to SPF list Dec 2 16:08:04 OSPF add LSA Router 10.10.134.11 distance 1 to SPF list Dec 2 16:08:04 IP nexthop so-1/1/0.0 0.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.10.134.12 distance 1 to SPF list Dec 2 16:08:04 IP nexthop so-1/1/1.0 0.0.0.0

Analyze OSPF Link-State Advertisement Packets in Detail

Action

To analyze OSPF link-state advertisement packets in detail, follow these steps:

1. In configuration mode, go to the following hierarchy level:


```
[edit]
user@host# edit protocols ospf traceoptions
```

2. Configure OSPF link-state packages:

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-update detail
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit protocols ospf traceoptions]
user@host# show
file ospf size 5m world-readable;
flag hello detail;
flag lsa-update detail;
```

4. Commit the configuration:

```
user@host# commit
```

5. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
user@host# run show log ospf
```

```
Dec 2 16:23:47 OSPF sent LSUpdate (4) -> 224.0.0.5 (so-1/1/0.0) ec 2 16:23:47
Version 2, length 196, ID 10.0.0.6, area 0.0.0.0
Dec 2 16:23:47 checksum 0xcc46, authtype 0
Dec 2 16:23:47 adv count 6 Dec 2 16:23:47 OSPF sent LSUpdate (4) -> 224.0.0.5
(so-1/1/1.0)
Dec 2 16:23:47 Version 2, length 196, ID 10.0.0.6, area 0.0.0.0 Dec 2
16:23:47 checksum 0xcc46, authtype 0
Dec 2 16:23:47 adv count 6
```


Diagnose OSPF Session Establishment Problems

Action

To trace OSPF messages in detail, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols ospf traceoptions
```

2. Configure OSPF hello messages:

```
[edit protocols ospf traceoptions]
user@host# set flag hello detail
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit protocols ospf traceoptions]
user@host# show
file ospf size 5m world-readable;
flag hello detail;
```

4. Commit the configuration:

```
user@host# commit
```

5. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
user@host# run show log ospf
```

```
Dec 2 16:14:24 Version 2, length 44, ID 10.0.0.6, area 1.0.0.0
Dec 2 16:14:24 checksum 0xf01a, authtype 0
Dec 2 16:14:24 mask 0.0.0.0, hello_ivl 10, opts 0x2, prio 128
Dec 2 16:14:24 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
Dec 2 16:14:24 OSPF sent Hello (1) -> 224.0.0.5 (so-1/1/2.0)
```



```

Dec 2 16:14:24 Version 2, length 44, ID 10.0.0.6, area 1.0.0.0
Dec 2 16:14:24 checksum 0xf01a, authtype 0
Dec 2 16:14:24 mask 0.0.0.0, hello_ivl 10, opts 0x2, prio 128
Dec 2 16:14:24 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
Dec 2 16:14:26 OSPF rcvd Hello 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0)
Dec 2 16:14:26 Version 2, length 48, ID 10.10.134.12, area 0.0.0.0
Dec 2 16:14:26 checksum 0x99b8, authtype 0Dec 2 16:14:26 mask
255.255.255.252, hello_ivl 10, opts 0x2, prio 1
Dec 2 16:14:26 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0
Dec 2 16:14:29 OSPF rcvd Hello 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0)
Dec 2 16:14:29 Version 2, length 48, ID 10.108.134.11, area 0.0.0.0
Dec 2 16:14:29 checksum 0x99b9, authtype 0Dec 2 16:14:29 mask
255.255.255.252, hello_ivl 10, opts 0x2, prio 1
Dec 2 16:14:29 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0

```

Meaning

[Table 77 on page 2506](#) lists OSPF tracing flags and presents example output for some of the flags.

Table 78: OSPF Protocol Tracing Flags

Tracing Flags	Description	Example Output
database description	All database description packets	<pre> Dec 2 15:44:51 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.29 (so-1/1/0.0) state changed from Full to Down Dec 2 15:44:51 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Full to Down Dec 2 15:44:55 RPD_OSPF_NBRUP: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Init to ExStart Dec 2 15:44:55 OSPF sent DbD (2) -> 224.0.0.5 (so-1/1/1.0) Dec 2 15:44:55 Version 2, length 32, ID 10.0.0.6, area 0.0.0.0 Dec 2 15:44:55 checksum 0xf76b, authtype 0 Dec 2 15:44:55 options 0x42, i 1, m 1, ms 1, seq 0xa009eee, mtu 4470 Dec 2 15:44:55 OSPF rcvd DbD 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0) Dec 2 15:44:55 Version 2, length 32, ID 10.10.134.12, area 0.0.0.0 Dec 2 15:44:55 checksum 0x312c, authtype 0 Dec 2 15:44:55 options 0x42, i 1, m 1, ms 1, seq 0x2154, mtu 4470 </pre>

Table 78: OSPF Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
error	OSPF errored packets	Dec 2 15:49:34 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:49:44 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:49:54 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:50:04 OSPF packet ignored: no matching interface from 172.16.120.29 Dec 2 15:50:14 OSPF packet ignored: no matching interface from 172.16.120.29
event	OSPF state transitions	Dec 2 15:52:35 OSPF interface ge-2/2/0.0 state changed from DR to DR Dec 2 15:52:35 OSPF interface ge-3/1/0.0 state changed from DR to DR Dec 2 15:52:35 OSPF interface ge-3/2/0.0 state changed from DR to DR Dec 2 15:52:35 OSPF interface ge-4/2/0.0 state changed from DR to DR Dec 2 15:53:21 OSPF neighbor 10.10.10.29 (so-1/1/0.0) state changed from Full to Down Dec 2 15:53:21 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.29 (so-1/1/0.0) state changed from Full to Down Dec 2 15:53:21 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Full to Down Dec 2 15:53:21 RPD_OSPF_NBRDOWN: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Full to Down Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Down to Init Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Init to ExStart Dec 2 15:53:25 RPD_OSPF_NBRUP: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Init to ExStart Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from ExStart to Exchange Dec 2 15:53:25 OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Exchange to Full Dec 2 15:53:25 RPD_OSPF_NBRUP: OSPF neighbor 10.10.10.33 (so-1/1/1.0) state changed from Exchange to Full

Table 78: OSPF Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
flooding	Link-state flooding packets	<p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 flooding on so-1/1/0.0</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 flooding on so-1/1/1.0</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 on no so-1/1/2.0 rexmit lists, no flood</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.218.0.0 10.0.0.6 on no so-1/1/3.0 rexmit lists, no flood</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.245.0.1 10.0.0.6 on no so-1/1/2.0 rexmit lists, no flood</p> <p>Dec 2 15:55:21 OSPF LSA Summary 10.245.0.1 10.0.0.6 on no so-1/1/3.0 rexmit lists, no flood</p>
hello	Hello packets	<p>Dec 2 15:57:25 OSPF sent Hello (1) -> 224.0.0.5 (ge-3/1/0.0)</p> <p>Dec 2 15:57:25 Version 2, length 44, ID 10.0.0.6, area 2.0.0.0</p> <p>Dec 2 15:57:25 checksum 0xe43f, authtype 0</p> <p>Dec 2 15:57:25 mask 255.255.0.0, hello_ivl 10, opts 0x2, prio 128</p> <p>Dec 2 15:57:25 dead_ivl 40, DR 10.218.0.1, BDR 0.0.0.0</p> <p>Dec 2 15:57:25 OSPF rcvd Hello 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0)</p> <p>Dec 2 15:57:25 Version 2, length 48, ID 10.10.134.12, area 0.0.0.0</p> <p>Dec 2 15:57:25 checksum 0x99b8, authtype 0</p> <p>Dec 2 15:57:25 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 1</p> <p>Dec 2 15:57:25 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0</p> <p>Dec 2 15:57:27 OSPF sent Hello (1) -> 224.0.0.5 (ge-3/2/0.0)</p> <p>Dec 2 15:57:27 Version 2, length 44, ID 10.0.0.6, area 2.0.0.0</p> <p>Dec 2 15:57:27 checksum 0xe4a5, authtype 0</p> <p>Dec 2 15:57:27 mask 255.255.0.0, hello_ivl 10, opts 0x2, prio 128</p> <p>Dec 2 15:57:27 dead_ivl 40, DR 10.116.0.1, BDR 0.0.0.0</p> <p>Dec 2 15:57:28 OSPF rcvd Hello 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0)</p> <p>Dec 2 15:57:28 Version 2, length 48, ID 10.10.134.11, area 0.0.0.0</p> <p>Dec 2 15:57:28 checksum 0x99b9, authtype 0</p> <p>Dec 2 15:57:28 mask 255.255.255.252, hello_ivl 10, opts 0x2, prio 1</p> <p>Dec 2 15:57:28 dead_ivl 40, DR 0.0.0.0, BDR 0.0.0.0</p>

Table 78: OSPF Protocol Tracing Flags (continued)

Tracing Flags	Description	Example Output
lsa-ack	Link-state acknowledgment packets	Dec 2 16:00:11 OSPF rcvd LSAck 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0) Dec 2 16:00:11 Version 2, length 44, ID 10.10.134.11, area 0.0.0.0 Dec 2 16:00:11 checksum 0xcdbf, authtype 0 Dec 2 16:00:11 OSPF rcvd LSAck 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0) Dec 2 16:00:11 Version 2, length 144, ID 10.10.134.12, area 0.0.0.0 Dec 2 16:00:11 checksum 0x73bc, authtype 0 Dec 2 16:00:16 OSPF rcvd LSAck 10.10.10.33 -> 224.0.0.5 (so-1/1/1.0) Dec 2 16:00:16 Version 2, length 44, ID 10.10.134.12, area 0.0.0.0 Dec 2 16:00:16 checksum 0x8180, authtype 0
lsa-request	Link-state request packets	Dec 2 16:01:38 OSPF rcvd LSReq 10.10.10.29 -> 224.0.0.5 (so-1/1/0.0) Dec 2 16:01:38 Version 2, length 108, ID 10.10.134.11, area 0.0.0.0 Dec 2 16:01:38 checksum 0xe86, authtype 0
lsa-update	Link-state update packets	Dec 2 16:09:12 OSPF built router LSA, area 0.0.0.0 Dec 2 16:09:12 OSPF built router LSA, area 1.0.0.0 Dec 2 16:09:12 OSPF built router LSA, area 2.0.0.0 Dec 2 16:09:13 OSPF sent LSUpdate (4) -> 224.0.0.5 (so-1/1/0.0) Dec 2 16:09:13 Version 2, length 268, ID 10.0.0.6, area 0.0.0.0 Dec 2 16:09:13 checksum 0x8047, authtype 0 Dec 2 16:09:13 adv count 7 Dec 2 16:09:13 OSPF sent LSUpdate (4) -> 224.0.0.5 (so-1/1/1.0) Dec 2 16:09:13 Version 2, length 268, ID 10.0.0.6, area 0.0.0.0 Dec 2 16:09:13 checksum 0x8047, authtype 0 Dec 2 16:09:13 adv count 7
packets	All OSPF packets	Not available.
packet-dump	Dump the contents of selected packet types	Not available.

Table 78: OSPF Protocol Tracing Flags *(continued)*

Tracing Flags	Description	Example Output
spf	SPF calculations	Dec 2 16:08:03 OSPF full SPF refresh scheduled Dec 2 16:08:04 OSPF SPF start, area 1.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.0.0.6 distance 0 to SPF list Dec 2 16:08:04 SPF elapsed time 0.000525s Dec 2 16:08:04 Stub elapsed time 0.000263s Dec 2 16:08:04 OSPF SPF start, area 2.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.0.0.6 distance 0 to SPF list Dec 2 16:08:04 SPF elapsed time 0.000253s Dec 2 16:08:04 Stub elapsed time 0.000249s Dec 2 16:08:04 OSPF SPF start, area 0.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.0.0.6 distance 0 to SPF list Dec 2 16:08:04 OSPF add LSA Router 10.10.134.11 distance 1 to SPF list Dec 2 16:08:04 IP nexthop so-1/1/0.0 0.0.0.0 Dec 2 16:08:04 OSPF add LSA Router 10.10.134.12 distance 1 to SPF list Dec 2 16:08:04 IP nexthop so-1/1/1.0 0.0.0.0

Analyze OSPF Link-State Advertisement Packets in Detail

Action

To analyze OSPF link-state advertisement packets in detail, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocols ospf traceoptions
```

2. Configure OSPF link-state packages:

```
[edit protocols ospf traceoptions]
user@host# set flag lsa-update detail
```

3. Verify the configuration:

```
user@host# show
```

For example:


```
[edit protocols ospf traceoptions]
user@host# show
file ospf size 5m world-readable;
flag hello detail;
flag lsa-update detail;
```

4. Commit the configuration:

```
user@host# commit
```

5. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
user@host# run show log ospf
```

```
Dec 2 16:23:47 OSPF sent LSUpdate (4) -> 224.0.0.5 (so-1/1/0.0) ec 2 16:23:47
Version 2, length 196, ID 10.0.0.6, area 0.0.0.0
Dec 2 16:23:47 checksum 0xcc46, authtype 0
Dec 2 16:23:47 adv count 6 Dec 2 16:23:47 OSPF sent LSUpdate (4) -> 224.0.0.5
(so-1/1/1.0)
Dec 2 16:23:47 Version 2, length 196, ID 10.0.0.6, area 0.0.0.0 Dec 2
16:23:47 checksum 0xcc46, authtype 0
Dec 2 16:23:47 adv count 6
```

Chassis

Chassis dimensions are listed in the physical specifications table for each router. For more information about chassis dimensions, see the router platform-specific hardware guide.

Each Juniper Networks router features a rigid sheet metal chassis that houses all of the router components. The chassis are designed to install into a variety of racks, including standard 19-inch equipment racks, telco center-mount racks, and four-post racks and cabinets. See [Table 79 on page 2518](#) for the maximum number of each router type that can be installed into a rack. Each chassis includes mounting ears or support posts to facilitate rack mounting, and one or more points for connecting an electrostatic discharge (ESD) wrist strap for use when servicing the router.

Table 79: Maximum Number of Routers per Rack

Router or Routing Node	Maximum in Standard Rack
T640	2
T320	3
M160	2
M40e	2
M40	2
M20	5
M5 and M10	14

Each chassis includes a midplane (called the backplane on an M40 router). The midplane transfers data packets to and from the FPCs, distributes power to router components, and provides signal connectivity to the router components for system monitoring and control.

Physical Interface Cards

Juniper Networks M-series routers and T-series platforms use PICs to connect to a wide variety of network media. PICs receive incoming packets from the network and transmit outgoing packets to the network, performing framing and line-speed signaling for their specific media type. Before transmitting outgoing data packets, the PICs encapsulate the packets received from the FPCs. Each PIC is equipped with an ASIC that performs control functions specific to the PIC's media type. For information about PICs, see the specific PIC guide.

Routing Engine

The Routing Engine consists of an Intel-based PCI platform running the Junos OS. The Routing Engine maintains the routing tables used by the router in which it is installed and controls the routing protocols on the router. The T640 routing node, and the T320, M160, M40e, and M20 routers support up to two Routing Engines, while the M40, M10, and M5 routers support a single Routing Engine.

Each Routing Engine consists of a CPU; SDRAM for storage of the routing and forwarding tables and other processes; a compact flash disk for primary storage of software images, configuration files, and microcode;

a hard disk for secondary storage; a PC card slot (on some M40 routers, a floppy disk) for storage of software upgrades; and interfaces for out-of-band management access.

Compare Information Logged Before and After the Reinstall

Purpose

Compare the operation of the system before and after the reinstall to ensure that everything is working as expected.

Action

To obtain system information, use the following commands:

```
user@host> show version
user@host> show chassis hardware
user@host> show chassis environment
user@host> show system boot-messages
user@host> show configuration
user@host> show interface terse
user@host> show bgp summary
user@host> show isis adjacency brief
user@host> show ospf neighbor brief
user@host> show system storage
```

Compare the information from these commands with the information you obtained before the reinstall.

Back Up the New Software

Purpose

After a week or so, when you are satisfied that the new software is running successfully, we recommend that you back up the reinstalled software.

Action

To back up the reinstalled software, use the following Junos OS CLI operational mode command:

```
user@host> request system snapshot
```

The root file system is backed up to **/altroot**, and **/config** is backed up to **/altconfig**. The root and **/config** file systems are on the router's internal flash drive, and the **/altroot** and **/altconfig** file systems are on the router's hard drive.

NOTE: After you issue the **request system snapshot** command, you cannot return to the previous version of the software because the running and backup copies of the software are identical.

Monitor Hardware Components

Purpose

NOTE: If the System Control Board (SCB), System and Switch Board (SSB), or Forwarding Engine Board (FEB) is not running, information about chassis components is not available through the command-line interface (CLI).

Action

To use the CLI to monitor Juniper Networks routers, follow these steps:

1. Log in to the router. The CLI operational mode prompt (>) appears.

If the operational mode prompt does not appear when you log in to the router, type **cli** to start the Junos OS and enter operational mode. The prompt changes to **>**, indicating that you are in operational mode.

2. Use one of the operational mode CLI commands listed in *Operational Mode CLI Commands for Router Monitoring* to monitor router hardware.

Log Software Version Information

Purpose

To log the Junos OS version information.

Action

To log the Junos OS version information, use the following Junos OS CLI operational mode command:

```
user@host> show version
```

Sample Output

```
user@host> show version
```



```

Hostname: host
Model: m10
JUNOS Base OS boot [5.0R5]
JUNOS Base OS Software Suite [5.0R5]
JUNOS Kernel Software Suite [5.0R5]
JUNOS Routing Software Suite [5.0R5]
JUNOS Packet Forwarding Engine Support [5.0R5]
JUNOS Crypto Software Suite [5.0R5]
JUNOS Online Documentation [5.0R5]
KERNEL 5.0R5 #0 built by builder on 2002-03-02 05:10:28 UTC
MGD release 5.0R5 built by builder on 2002-03-02 04:45:32 UTC
CLI release 5.0R5 built by builder on 2002-03-02 04:44:22 UTC
CHASSISD release 5.0R5 built by builder on 2002-03-02 04:43:37 UTC
DCD release 5.0R5 built by builder on 2002-03-02 04:42:47 UTC
RPD release 5.0R5 built by builder on 2002-03-02 04:46:17 UTC
SNMPD release 5.0R5 built by builder on 2002-03-02 04:52:26 UTC
MIB2D release 5.0R5 built by builder on 2002-03-02 04:45:37 UTC
APSD release 5.0R5 built by builder on 2002-03-02 04:43:31 UTC
VRRPD release 5.0R5 built by builder on 2002-03-02 04:52:34 UTC
ALARMD release 5.0R5 built by builder on 2002-03-02 04:43:24 UTC
PFED release 5.0R5 built by builder on 2002-03-02 04:46:06 UTC
CRAFTD release 5.0R5 built by builder on 2002-03-02 04:44:30 UTC
SAMPLED release 5.0R5 built by builder on 2002-03-02 04:52:20 UTC
ILMID release 5.0R5 built by builder on 2002-03-02 04:45:21 UTC
BPRELAYD release 5.0R5 built by builder on 2002-03-02 04:42:41 UTC
RMOPD release 5.0R5 built by builder on 2002-03-02 04:46:11 UTC
jkernel-dd release 5.0R5 built by builder on 2002-03-02 04:41:07 UTC
jroute-dd release 5.0R5 built by builder on 2002-03-02 04:41:21 UTC
jdocs-dd release 5.0R5 built by builder on 2002-03-02 04:39:11 UTC

```

Meaning

The sample output shows the hostname, router model, and the different Junos OS packages, processes, and documents.

Hardware Components

IN THIS SECTION

- [Chassis | 2522](#)
- [Flexible PIC Concentrators | 2523](#)

- Physical Interface Cards | 2523
- Routing Engine | 2523
- Power Supplies | 2523
- Cooling System | 2524

Each Juniper Networks router consists of a chassis and a set of components, including FPCs, PICs, Routing Engines, power supplies, cooling system, and cable management system. Many of the components are field-replaceable units. The following major components are discussed in this topic:

Chassis

Chassis dimensions are listed in the physical specifications table for each router. For more information about chassis dimensions, see the router platform-specific hardware guide.

Each Juniper Networks router features a rigid sheet metal chassis that houses all of the router components. The chassis are designed to install into a variety of racks, including standard 19-inch equipment racks, telco center-mount racks, and four-post racks and cabinets. See [Table 79 on page 2518](#) for the maximum number of each router type that can be installed into a rack. Each chassis includes mounting ears or support posts to facilitate rack mounting, and one or more points for connecting an electrostatic discharge (ESD) wrist strap for use when servicing the router.

Table 80: Maximum Number of Routers per Rack

Router or Routing Node	Maximum in Standard Rack
T640	2
T320	3
M160	2
M40e	2
M40	2
M20	5
M5 and M10	14

Each chassis includes a midplane (called the backplane on an M40 router). The midplane transfers data packets to and from the FPCs, distributes power to router components, and provides signal connectivity to the router components for system monitoring and control.

Flexible PIC Concentrators

The FPCs house the PICs used in the router and connect them to other router components. FPCs install into the front of the router in either a vertical or horizontal orientation, depending on the router. A compatible FPC can be installed into any available FPC slot, regardless of the PICs it contains. If a slot is not occupied by an FPC, a blank FPC panel must be installed to shield the empty slot and allow cooling air to circulate properly through the FPC card cage. For information about FPCs, see the specific hardware guide.

Physical Interface Cards

Juniper Networks M-series routers and T-series platforms use PICs to connect to a wide variety of network media. PICs receive incoming packets from the network and transmit outgoing packets to the network, performing framing and line-speed signaling for their specific media type. Before transmitting outgoing data packets, the PICs encapsulate the packets received from the FPCs. Each PIC is equipped with an ASIC that performs control functions specific to the PIC's media type. For information about PICs, see the specific PIC guide.

Routing Engine

The Routing Engine consists of an Intel-based PCI platform running the Junos OS. The Routing Engine maintains the routing tables used by the router in which it is installed and controls the routing protocols on the router. The T640 routing node, and the T320, M160, M40e, and M20 routers support up to two Routing Engines, while the M40, M10, and M5 routers support a single Routing Engine.

Each Routing Engine consists of a CPU; SDRAM for storage of the routing and forwarding tables and other processes; a compact flash disk for primary storage of software images, configuration files, and microcode; a hard disk for secondary storage; a PC card slot (on some M40 routers, a floppy disk) for storage of software upgrades; and interfaces for out-of-band management access.

Power Supplies

Each Juniper Networks M-series router, T-series platform, or MX-series router has one to four load-sharing power supplies depending on the platform. If a power supply in a redundant configuration is removed or fails, the other power supplies assume the electrical load. For more information about the power supplies in each router, see the router platform-specific hardware guide.

The power supplies are connected to the router midplane, which distributes the different output voltages throughout the router and its components. Some routers can operate using either AC or DC power; other

routers operate with DC power only. For information about the type of power used by each router, see the platform-specific hardware guide.

Cooling System

Each Juniper Networks M-series router and T-series platform features a cooling system designed to keep all router components within recommended operating temperature limits. If one component of the cooling system fails or is removed, the system automatically adjusts the speed of the remaining components to keep the temperature within the acceptable range. The cooling system for each router is unique and can consist of fans, impellers, and air filters. For information about the cooling system components of each router, see the “Major Hardware Components” table in the router platform-specific hardware guide.

10

PART

Configuration Statements

MPLS Configuration Statements | **2527**

RSVP Configuration Statements | **2829**

LDP Configuration Statements | **2905**

CCC and TCC Configuration Statements | **2999**

GMPLS Configuration Statements | **3021**

PCEP Configuration Statements | **3071**

MPLS Configuration Statements

IN THIS CHAPTER

- [abstract-hop](#) | 2534
- [adaptive](#) | 2536
- [adjust-interval](#) | 2537
- [adjust-threshold](#) | 2538
- [adjust-threshold-absolute](#) | 2539
- [adjust-threshold-activate-bandwidth](#) | 2540
- [adjust-threshold-overflow-limit](#) | 2541
- [adjust-threshold-underflow-limit](#) | 2542
- [admin-down](#) | 2543
- [admin-group \(for Interfaces\)](#) | 2544
- [admin-group \(for LSPs\)](#) | 2545
- [admin-group-extended](#) | 2546
- [admin-groups](#) | 2548
- [admin-groups-extended](#) | 2549
- [admin-groups-extended-range](#) | 2551
- [advertise-mode \(MPLS\)](#) | 2553
- [advertisement-hold-time](#) | 2554
- [allow-fragmentation](#) | 2555
- [always-mark-connection-protection-tlv](#) | 2556
- [associate-backup-pe-groups](#) | 2557
- [associate-lsp](#) | 2558
- [auto-bandwidth \(MPLS Tunnel\)](#) | 2559
- [auto-bandwidth \(MPLS Statistics\)](#) | 2561
- [auto-policing](#) | 2562
- [backup-pe-group](#) | 2563
- [bandwidth \(Fast Reroute, Signaled, and Multiclass LSPs\)](#) | 2564
- [bandwidth \(Static LSP\)](#) | 2566
- [bandwidth-model](#) | 2567

- bandwidth-percent | 2568
- bfd-liveness-detection (Protocols MPLS) | 2569
- bfd-liveness-detection (LSP) | 2571
- class-of-service (Protocols MPLS) | 2573
- compute-options | 2574
- compute-profile | 2575
- connections (MPLS) | 2577
- constituent-list | 2578
- container-label-switched-path | 2579
- corouted-bidirectional | 2580
- corouted-bidirectional-passive | 2581
- credibility | 2582
- database | 2584
- delay (querier) | 2585
- delay (responder) | 2586
- description (Protocols MPLS) | 2587
- description (Protocols Layer 2 VPN) | 2588
- deselect-on-bandwidth-failure | 2589
- diffserv-te | 2590
- disable (Protocols MPLS) | 2591
- dual-transport | 2592
- dynamic (Source Packet Routing) | 2593
- dynamic-tunnels | 2595
- egress-protection (MPLS) | 2597
- encapsulation-type (Layer 2 VPNs) | 2599
- encoding-type | 2601
- entropy-label | 2602
- entropy-label | 2603
- ethernet-vlan (Protocols Link Management) | 2604
- ether-pseudowire | 2605
- exclude (for Administrative Groups) | 2606
- exclude (for Fast Reroute) | 2607
- exclude-srlg | 2608
- exp | 2609

- [expand-loose-hop](#) | **2610**
- [explicit-null \(Protocols MPLS\)](#) | **2611**
- [export \(MPLS Traffic engineering database\)](#) | **2612**
- [failure-action \(Protocols MPLS\)](#) | **2613**
- [family](#) | **2614**
- [family mpls](#) | **2615**
- [fast-reroute \(Protocols MPLS\)](#) | **2619**
- [fate-sharing](#) | **2620**
- [forwarding-rib](#) | **2622**
- [forwarding-table](#) | **2623**
- [from \(Protocols MPLS\)](#) | **2624**
- [gpipid](#) | **2625**
- [gre \(Routing Options\)](#) | **2627**
- [hop-limit](#) | **2628**
- [import \(MPLS Traffic Engineering Database\)](#) | **2630**
- [ip-tunnel-rpf-check](#) | **2632**
- [ipv4 \(Family MPLS\)](#) | **2634**
- [ipv6 \(Family MPLS\)](#) | **2636**
- [ip-version \(Family MPLS\)](#) | **2637**
- [include-all \(for Administrative Groups\)](#) | **2638**
- [include-all \(for Fast Reroute\)](#) | **2639**
- [include-any \(for Administrative Groups\)](#) | **2640**
- [include-any \(for Fast Reroute\)](#) | **2641**
- [ingress \(LSP\)](#) | **2642**
- [install \(Protocols MPLS\)](#) | **2644**
- [ingress-policy](#) | **2647**
- [interface \(Protocols MPLS\)](#) | **2648**
- [interface \(MPLS\)](#) | **2649**
- [inter-domain](#) | **2650**
- [ip-tunnel-rpf-check](#) | **2651**
- [ipv6-tunneling](#) | **2652**
- [label-switched-path \(Protocols MPLS\)](#) | **2653**
- [label-switched-path](#) | **2659**
- [label-switched-path-template \(Container LSP\)](#) | **2660**

- [ldp-tunneling](#) | 2661
- [least-fill](#) | 2661
- [link-protection \(Dynamic LSPs\)](#) | 2662
- [link-protection \(Static LSPs\)](#) | 2663
- [load-balance-label-capability](#) | 2664
- [log-updown \(Protocols MPLS\)](#) | 2665
- [longest-match](#) | 2667
- [loss \(querier\)](#) | 2668
- [loss \(responder\)](#) | 2669
- [loss-delay \(querier\)](#) | 2670
- [lsp-attributes](#) | 2671
- [lsping-channel-type](#) | 2672
- [l2vpn](#) | 2673
- [maximum-bandwidth \(Protocols MPLS\)](#) | 2676
- [maximum-helper-recovery-time](#) | 2677
- [maximum-helper-restart-time \(RSVP\)](#) | 2678
- [maximum-labels](#) | 2679
- [minimum-bandwidth-adjust-interval](#) | 2680
- [minimum-bandwidth-adjust-threshold-change](#) | 2681
- [minimum-bandwidth-adjust-threshold-value](#) | 2682
- [metric \(Protocols MPLS\)](#) | 2683
- [minimum-bandwidth](#) | 2684
- [monitor-bandwidth](#) | 2685
- [most-fill](#) | 2685
- [mpls \(Protocols\)](#) | 2686
- [mpls](#) | 2687
- [mpls-tp-mode](#) | 2690
- [mtu-signaling](#) | 2691
- [neighbor \(Protocols Layer 2 Circuit\)](#) | 2692
- [next-hop \(Protocols MPLS\)](#) | 2694
- [no-bfd-triggered-local-repair](#) | 2695
- [no-cspf](#) | 2696
- [no-decrement-ttl](#) | 2698
- [graceful-restart \(Enabling Globally\)](#) | 2699

- [helper-disable \(Multiple Protocols\) | 2701](#)
- [no-install-to-address | 2702](#)
- [no-load-balance-label-capability | 2703](#)
- [no-mcast-replication | 2704](#)
- [no-propagate-ttl | 2705](#)
- [no-transit-statistics | 2706](#)
- [no-trap | 2707](#)
- [node-protection \(Static LSP\) | 2708](#)
- [normalization | 2709](#)
- [oam \(Protocols MPLS\) | 2711](#)
- [optimize-adaptive-teardown | 2713](#)
- [optimize-aggressive | 2714](#)
- [optimize-hold-dead-delay | 2715](#)
- [optimize-switchover-delay | 2716](#)
- [optimize-timer \(Protocols MPLS\) | 2717](#)
- [p2mp \(Protocols MPLS\) | 2719](#)
- [p2mp-lsp-next-hop | 2720](#)
- [path \(Protocols MPLS\) | 2721](#)
- [path | 2723](#)
- [path-mtu | 2724](#)
- [per-prefix-label | 2725](#)
- [performance-monitoring \(Protocols MPLS\) | 2727](#)
- [policing \(Protocols MPLS\) | 2729](#)
- [policing | 2730](#)
- [policy-multipath | 2731](#)
- [policy-statement | 2733](#)
- [pop | 2739](#)
- [pop-and-forward \(Protocols MPLS\) | 2740](#)
- [preference \(Protocols MPLS\) | 2741](#)
- [primary \(Protocols MPLS\) | 2743](#)
- [primary | 2744](#)
- [priority \(Protocols MPLS\) | 2745](#)
- [protection-revert-time | 2746](#)
- [push | 2747](#)

- random | 2748
- record | 2749
- remote-interface-switch | 2750
- remote-site-id | 2752
- retry-limit | 2753
- retry-timer | 2754
- revert-timer | 2755
- revert-timer | 2756
- resignal-minimum-bandwidth | 2757
- resolution-map | 2758
- responder (performance-monitoring) | 2759
- rpf-check-policy (Routing Options) | 2760
- rsvp-error-hold-time | 2761
- sampling (Protocols MPLS) | 2762
- sbfd | 2764
- secondary (Protocols MPLS) | 2765
- secondary | 2767
- segment | 2768
- segment-list | 2770
- select | 2775
- signal-bandwidth | 2776
- signaling | 2777
- site (Layer 2 Circuits) | 2778
- site-identifier (Layer 2 Circuits) | 2779
- smart-optimize-timer | 2780
- soft-preemption (Protocols MPLS) | 2781
- source-routing-path | 2782
- source-routing-path-template | 2785
- source-routing-path-template-map | 2787
- splitting-merging | 2788
- spring-te (Dynamic Tunnels) | 2790
- srgb-label-range | 2791
- srlg | 2792
- srlg-cost | 2793

- srlg-value | 2794
- standby | 2795
- standby | 2796
- static-label-switched-path | 2797
- statistics (Protocols MPLS) | 2799
- swap | 2801
- switch-away-lsps | 2802
- switching-type | 2803
- sync-active-path-bandwidth | 2804
- te-class-matrix | 2806
- to | 2808
- traceoptions (Protocols MPLS) | 2809
- traffic-class (delay) | 2812
- traffic-class (loss) | 2814
- traffic-class (loss-delay) | 2816
- traffic-engineering (Protocols MPLS) | 2818
- traffic-engineering | 2820
- traffic-engineering (Protocols BGP) | 2821
- transit-lsp-association | 2823
- ultimate-hop-popping | 2824
- vrf-table-label | 2826

abstract-hop

Syntax

```
abstract-hop abstract-hop-name {
    constituent-list constituent-list-name (include-any-list | include-all-list | exclude-all-list | exclude-any-list);
    operators (AND | OR);
}
```

Hierarchy Level

```
[edit logical systems logical-systems-name protocols mpls]
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 17.1 for all platforms.

Description

Define router clusters or groups, similar to the sequence of real-hop constraints (strict or loose), as a sequence of abstract hops for setting up a label-switched path (LSP).

An abstract hop is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs), along with the ordering property of real hops. As a result, when a sequence of abstract hops is used in a path constraint, ordering is achieved among the groups of routers that meet a logical combination of link or node attributes called constituent attributes. A path can use a combination of real and abstract hops as constraints.

Options

abstract-hop-name—Name of the abstract hop that is a logical combination of the existing traffic engineering constraints, such as administrative groups, extended administrative groups, and SRLGs, along with the ordering property of real hops.

constituent-list constituent-list-name—Name of the predefined constituent list to be included in defining the abstract hop. A constituent list enables you to define a set of constituent attributes that is identified with a user-defined name.

include-any-list—Satisfy any one of the attributes specified in the constituent list.

include-all-list—Satisfy all of the attributes specified in the constituent list.

exclude-all-list—Satisfy none of the attributes specified in the constituent list.

exclude-any-list—Fail to satisfy any one of the attributes specified in the constituent list.

operators—Specify the operation between constituent lists when more than one constituent list is included in the abstract hop definition.

AND—Satisfy all the constituent lists referenced in the abstract hop definition for the attached node to be a member of the abstract hop.

OR—Satisfy at least one of the constituent lists referenced in the abstract hop definition for the attached node to be a member of the abstract hop.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring Abstract Hops for MPLS LSPs | 446](#)

[constituent-list | 2578](#)

[show mpls abstract-hop-membership | 3170](#)

[show mpls lsp abstract-computation | 3240](#)

adaptive

Syntax

```
adaptive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |  
  secondary) path-name],  
[edit protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

During reroute, do not double-count bandwidth on links shared by the old and new paths. Including this statement causes RSVP to use shared explicit (SE) reservation styles and assists in smooth transition during rerouting.

Default

The configured object is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Adaptive LSP Configuration](#) | 564

adjust-interval

Syntax

```
adjust-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify the bandwidth reallocation interval.

Options

seconds—Bandwidth reallocation interval, in seconds.

Range: 300 through 315,360,000 seconds

Default: 86,400 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Automatic Bandwidth Allocation Interval](#) | 550

adjust-threshold

Syntax

```
adjust-threshold percent;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify in percentage how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.

To specify the changes in the automatic bandwidth adjustment for a LSP in absolute value, use the [adjust-threshold-absolute](#) statement instead.

Options

percent—Bandwidth demand for the current bandwidth adjustment interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the percentage specified by this statement, the LSP's bandwidth is adjusted to the current bandwidth demand.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Automatic Bandwidth Adjustment Threshold](#) | 551

adjust-threshold-absolute

Syntax

```
adjust-threshold-absolute bps;
```

Hierarchy Level

```
[edit logical-systems name protocols mpls label-switched-path name auto-bandwidth ],
[edit logical-systems name routing-instances name protocols mpls label-switched-path name auto-bandwidth ],
[edit protocols mpls label-switched-path name auto-bandwidth ],
[edit routing-instances name protocols mpls label-switched-path name auto-bandwidth ]
```

Release Information

Statement introduced before Junos OS Release 17.4R1 on all platforms.

Description

Specify in bits per second how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in the average LSP utilization.

The **adjust-threshold-absolute** statement works in conjunction with the **adjust-threshold** statement, which specifies the change in automatic bandwidth adjustment for an LSP as a percentage.

By triggering automatic bandwidth LSP resignaling based on absolute change in bandwidth instead of percentage bandwidth change, LSP resignaling can be optimized for both big and small LSPs at the same time.

Options

bps—Change in average LSP utilization to trigger automatic bandwidth adjustment in bits per second.

Default: 0 bps

Required Privilege Level

routing

RELATED DOCUMENTATION

[Configuring the Automatic Bandwidth Adjustment Threshold](#) | 551

adjust-threshold-activate-bandwidth

Syntax

```
adjust-threshold-activate-bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced before Junos OS Release 14.1.

Description

Specify an absolute value to prevent automatic adjustment of signaled bandwidth and aggressive re-signaling of a label-switched path (LSP) when the actual bandwidth over the LSP is below the configured threshold, although the adjust-threshold percentage condition is satisfied.

Options

bps—Amount of bandwidth that is compared with the maximum of all traffic samples during an adjustment interval. If the maximum average bandwidth is less than this configured value, automatic bandwidth adjustment or re-signaling does not happen, even if the adjust-threshold percentage condition is satisfied.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Automatic Bandwidth Adjustment Threshold](#) | 551

adjust-threshold-overflow-limit

Syntax

```
adjust-threshold-overflow-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced in Junos OS Release 7.5.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.

Options

number—Number of consecutive bandwidth overflow samples.

Range: 1 through 65,535

Default: This feature is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring a Limit on Bandwidth Overflow and Underflow Samples](#) | 552

adjust-threshold-underflow-limit

Syntax

```
adjust-threshold-underflow-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced in Junos OS Release 11.3.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify the number of consecutive bandwidth underflow samples before triggering a bandwidth adjustment.

Options

number—Number of consecutive bandwidth underflow samples.

Range: 1 through 65,535

Default: This feature is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring a Limit on Bandwidth Overflow and Underflow Samples](#) | 552

admin-down

Syntax

```
admin-down;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 8.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Set a nonpacket GMPLS LSP to the administrative down state. This statement does not affect control path setup or data forwarding for packet LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS](#) | 1374

admin-group (for Interfaces)

Syntax

```
admin-group [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls interface interface-name],  
[edit protocols mpls interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define administrative groups for an interface.

Options

group-names—One or more names of groups defined with the [admin-groups](#) statement at the [\[edit protocols mpls\]](#) hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Administrative Groups for LSPs | 532](#)

[admin-groups | 2548](#)

admin-group (for LSPs)

Syntax

```
admin-group {  
  exclude [ group-names ];  
  include-all [ group-names ];  
  include-any [ group-names ];  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |  
  secondary) path-name],  
[edit protocols mpls],  
[edit protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define the administrative groups to include or exclude an LSP and a path's primary and secondary paths.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Administrative Groups for LSPs](#) | 532

admin-group-extended

Syntax

```
admin-group-extended {
  apply-groups group-value;
  apply-groups-except group-value;
  exclude [ group-values ];
  include-all [ group-values ];
  include-any [ group-values ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary | secondary)
  path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced in Junos OS Release 11.1.

Description

Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the **admin-groups-extended-range** statement. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.

Options

apply-groups—Apply the specified administrative groups for the LSP or for the primary and secondary paths.

apply-groups-except—Exclude the specified administrative groups from the LSP or from the primary and secondary paths.

exclude—Define the administrative groups to exclude from an LSP or from the primary and secondary paths.

include-all—Require the LSP to traverse links that include all of the defined administrative groups.

include-any—Define the administrative groups to include for an LSP for the primary and secondary paths.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

	Configuring Extended Administrative Groups for LSPs 534
	Configuring Administrative Groups for LSPs 532
	admin-groups-extended 2549
	admin-groups-extended-range 2551

admin-groups

Syntax

```
admin-groups {  
    group-name group-value;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure administrative groups to implement link coloring of resource classes.

Options

group-name—Name of the group. You can assign up to 32 names. The names and their corresponding values must be identical across all routers within a single domain.

group-value—Value assigned to the group. The names and their corresponding values must be identical across all routers within a single domain.

Range: 0 through 31

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Administrative Groups for LSPs | 532](#)

[admin-group \(for Interfaces\) | 2544](#)

admin-groups-extended

Syntax

```
admin-groups-extended group-name {
  group-value group-identifier;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls interface interface-name],
[edit logical-systems logical-system-name routing-options],
[edit protocols mpls interface interface-name],
[edit routing-options]
```

Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the **admin-groups-extended-range** statement. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.

Options

group-name—The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.

group-value group-identifier—The group identifier must be within the range of configurable values, 32 and 4,294,967,295.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Extended Administrative Groups for LSPs | 534](#)

[Configuring Administrative Groups for LSPs | 532](#)

admin-group-extended | 2546

admin-groups-extended-range | 2551

admin-groups-extended-range

Syntax

```
admin-groups-extended-range {
    maximum maximum-number;
    minimum minimum-number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit routing-options]
```

Release Information

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Enables you to configure extended administrative groups, represented by a 32-bit value, expanding the number of administrative groups supported in the network beyond just 32. In MPLS traffic engineering, a link can be configured with a set of administrative groups (also known as colors or resource classes). Administrative groups are carried in IGPs (OSPFv2 and IS-IS) as a 32-bit value assigned to each link. By default, Juniper Networks routers interpret this 32-bit value as a bit mask with each bit representing a group. This normally limits each network to a total of 32 distinct administrative groups (value range 0 through 31).

The extended administrative groups configuration accepts a set of interfaces with a corresponding set of extended administrative group names. It converts the names into a set of 32-bit values and propagates this information into the IGP. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.

Options

maximum *maximum-number*—The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.

minimum *minimum-number*—The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Extended Administrative Groups for LSPs 534
Configuring Administrative Groups for LSPs 532
admin-group-extended 2546

advertise-mode (MPLS)

Syntax

```
advertise-mode (stub-alias | stub-proxy);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls egress-protection context-identifier context-id],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name egress-protection
 context-identifier context-id],
[edit protocols mpls egress-protection context-identifier context-id],
[edit protocols mpls label-switched-path lsp-name egress-protection context-identifier context-id]
```

Release Information

Statement introduced in Junos OS Release 13.3.

Description

Configure the method for the interior gateway protocol (IGP) to advertise egress protection availability.

Egress protection availability is advertised in the IGP. Label protocols along with CSPF use this information to do egress protection.

Options

stub-alias—Context identifier has an alias.

In the alias method, the LSP end-point address has an explicit backup egress node where the backup node can be learned or configured on the penultimate hop node (PHN) of a protected LSP. With this model, the PHN of a protected LSP sets up the bypass LSP tunnel to back up the egress node by avoiding the primary egress node. This model requires a Junos OS upgrade in core nodes, but is flexible enough to support all traffic engineering constraints.

stub-proxy—Context-identifier has a stub proxy node.

A stub node is one that only appears at the end of an AS path, which means it does not provide transit service. In this mode, known as the virtual or proxy mode, the LSP end-point address is represented as a node with bidirectional links, with the LSP's primary egress node and backup egress node. With this representation, the penultimate hop of the LSP primary egress point can behave like a PLR in setting up a bypass tunnel to back up the egress by avoiding the primary egress node. This model has the advantage that you do not need to upgrade Junos OS on core nodes and thereby helps operators to deploy this technology.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Egress Protection for Layer 3 VPN Edge Protection Overview

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

advertisement-hold-time

Syntax

```
advertisement-hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Do not advertise when the LSP goes from up to down, for a certain period of time known as the hold time.

Options

seconds—Hold time, in seconds.

Range: 0 through 65,535 seconds

Default: 5 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Damping Advertisement of LSP State Changes | 565](#)

allow-fragmentation

Syntax

```
allow-fragmentation;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls path-mtu],  
[edit protocols mpls path-mtu]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Allow IP packets to be fragmented before they are encapsulated in MPLS.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Enabling Packet Fragmentation](#) | 854

always-mark-connection-protection-tlv

Syntax

```
always-mark-connection-protection-tlv;
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols mpls interface interface-name],  
[edit protocols mpls interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 10.2.

Description

(MX Series routers only) Enable you to switch an LSP away from a network node using a bypass LSP. This feature could be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

This statement marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, you then need to configure the **switch-away-lsps** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Switching LSPs Away from a Network Node](#) | 846

associate-backup-pe-groups

Syntax

```
associate-backup-pe-groups;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Description

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.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Enabling Point-to-Point LSPs to Monitor Egress PE Routers](#) | 727

associate-lsp

Syntax

```
associate-lsp lsp-name {  
    from from-ip-address;  
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name oam]
```

Release Information

Statement introduced in Junos OS Release 12.1.

Description

Configure associated bidirectional label-switched paths (LSPs) on the two ends of an LSP for sending and receiving GAL and G-Ach OAM messages.

Options

from *from-ip-address*—(Optional) Source address for the associated LSP configuration.

If omitted, this is derived from the **to** address of the ingress LSP configuration.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring the MPLS Transport Profile for OAM](#) | 1222

auto-bandwidth (MPLS Tunnel)

Syntax

```
auto-bandwidth {
  adjust-interval seconds;
  adjust-threshold percent;
  adjust-threshold-absolute;
  adjust-threshold-activate-bandwidth bps
  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;
  resignal-minimum-bandwidth
  sync-active-path-bandwidth
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for QFX Series switches.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.

NOTE: In calculating the value for **Max AvgBW** (relative to the ingress LSP), the sample collected during make before break (MBB) is ignored to prevent inaccurate results. The first sample after a bandwidth adjustment, or after a change in the LSP ID (regardless of path change), is also ignored.

Options

adjust-threshold-absolute *adjust-threshold-absolute-value*—Configure an absolute value based threshold along with the percentage based threshold that helps avoid the bandwidth getting triggered for LSPs of both small and large bandwidth reservations.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Automatic Bandwidth Allocation for LSPs | 547](#)

[request mpls lsp adjust-autobandwidth | 3142](#)

[show mpls lsp autobandwidth | 3243](#)

auto-bandwidth (MPLS Statistics)

Syntax

```
auto-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls statistics],  
[edit protocols mpls statistics]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for QFX Series switches.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Collect statistics related to automatic bandwidth.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Automatic Bandwidth Allocation for LSPs | 547](#)

[Configuring MPLS to Gather Statistics | 242](#)

[statistics | 2799](#)

auto-policing

Syntax

```
auto-policing {
  class all (drop | loss-priority-high | loss-priority-low);
  class ctnumber (drop | loss-priority-high | loss-priority-low);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for QFX10000 Series switches in release 15.1X53-D40.

Description

Enable the automatic policing of all the MPLS LSPs on the router or logical system.

Options

class all—Apply the same policer action to all the class types (**ct0**, **ct1**, **ct2**, and **ct3**).

class ctnumber—Specific class type (**ct0**, **ct1**, **ct2**, or **ct3**) to which to apply a policer action.

Policer actions—You can specify the following policer actions:

Default: no action

- **drop**—Drop all packets.
- **loss-priority-high**—Set the packet loss priority (PLP) to high.
- **loss-priority-low**—Set the PLP to low.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[policing \(Protocols MPLS\) | 2729](#)

[Configuring Automatic Policers | 124v](#)

backup-pe-group

Syntax

```
backup-pe-group group-name {  
    backups [ addresses ];  
    local-address address;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast],  
[edit logical-systems logical-system-name routing-options multicast],  
[edit routing-instances routing-instance-name routing-options multicast],  
[edit routing-options multicast]
```

Release Information

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options

group-name—Name of the group for PE backups.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Example: Configuring Ingress PE Redundancy*

bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)

Syntax

```
bandwidth bps {
  ct0 bps;
  ct1 bps;
  ct2 bps;
  ct3 bps;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
  secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name fast-reroute],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

When configuring an LSP, specify the traffic rate associated with the LSP.

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.

When configuring a multiclass LSP, use the **ctnumber bandwidth** statements to specify the bandwidth to be allocated for each class type.

Options

bps—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations **k** (for a thousand), **m** (for a million), or **g** (for a billion).

Range: Any positive integer

Default: 0 (no bandwidth is reserved)

NOTE: On the ACX Series, *bps* is the only supported option.

ctnumber bps—Bandwidth for the specified class type, in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations **k** (for a thousand), **m** (for a million), or **g** (for a billion [also called a thousand million]).

Range: Any positive integer

Default: 0 (no bandwidth is reserved)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute | 479](#)

[Configuring the Bandwidth Value for LSPs | 546](#)

[Configuring Traffic-Engineered LSPs | 1213](#)

[Configuring Class-Type Bandwidth Constraints for Multiclass LSPs | 1216](#)

bandwidth (Static LSP)

Syntax

```
bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name bypass],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],  
[edit protocols mpls static-label-switched-path lsp-name bypass],  
[edit protocols mpls static-label-switched-path lsp-name ingress],  
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Description

When configuring a static LSP, specify the traffic rate associated with the LSP.

Options

bps—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations **k** (for a thousand), **m** (for a million), or **g** (for a billion).

Range: Any positive integer

Default: 0 (no bandwidth is reserved)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Static LSPs](#) | 603

bandwidth-model

Syntax

```
bandwidth-model {  
    extended-mam;  
    mam;  
    rdm;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls diffserv-te],  
[edit protocols mpls diffserv-te]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Configure the bandwidth model for differentiated services. Note that you cannot configure both bandwidth models at the same time.

Options

extended-mam—The extended maximum allocation model (MAM) is a bandwidth model based on MAM.

mam—The MAM is defined in RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.

rdm—The Russian dolls bandwidth allocation model (RDM) is defined in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*. RDM makes efficient use of bandwidth by allowing the class types to share bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Bandwidth Model | 1201](#)

bandwidth-percent

Syntax

```
bandwidth-percent percentage;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],  
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the percentage of bandwidth to reserve for the detour path in case the primary path for a traffic engineered LSP or a multiclass LSP fails. The percentage configured indicates the percentage of the protected path's bandwidth that is reserved for the detour path.

Options

percentage—The percentage of the protected path's bandwidth that is reserved for the detour path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute | 479](#)

[Configuring Fast Reroute for Traffic-Engineered LSPs | 1214](#)

[Configuring Fast Reroute for Multiclass LSPs | 1218](#)

bfd-liveness-detection (Protocols MPLS)

Syntax

```

bfd-liveness-detection {
  failure-action {
    make-before-break teardown-timeout seconds;
    teardown;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  minimum-transmit-interval milliseconds;
  multiplier detection-time-multiplier;
}

```

Hierarchy Level

```

[edit protocols mpls label-switched-path lsp-name oam],
[edit protocols mpls oam]

```

Release Information

Statement introduced in Junos OS Release 7.6.

failure-action option added in Junos OS Release 8.5.

Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description

Enable Bidirectional Forwarding Detection (BFD) for all of the MPLS LSPs or for just a specific LSP.

Options

minimum-interval—Minimum transmit and receive interval.

Range: 50 through 255,000 milliseconds

Default: 50

minimum-receive-interval—Minimum receive interval.

Range: 50 through 255,000 milliseconds

Default: 50

minimum-transmit-interval—Minimum transmit interval.

Range: 50 through 255,000 milliseconds

Default: 50

multiplier—Detection time multiplier.

Range: 1 through 255

Default: 3

The **failure-action** statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring BFD for MPLS IPv4 LSPs | 109](#)

[Configuring Bidirectional Forwarding Detection for MPLS \(CLI Procedure\) | 101](#)

bfd-liveness-detection (LSP)

Syntax

```

bfd-liveness-detection {
  minimum-interval milliseconds;
  multiplier multiplier;
  no-router-alert-option;
  sbfd {
    remote-discriminator remote-discriminator;
  }
}

```

Hierarchy Level

[edit protocols source-packet-routing [segment-list](#)],

[edit protocols source-packet-routing source-routing-path *lsp-name* primary *segment-list-name*],

[edit protocols source-packet-routing source-routing-path *lsp-name* secondary *segment-list-name*]

Release Information

Statement introduced in Junos OS Release 18.4R1.

Support at the following hierarchy levels introduced in Junos OS Release 19.1R1: **[edit protocols source-packet-routing source-routing-path *lsp-name* primary *segment-list-name*]**, and **[edit protocols source-packet-routing source-routing-path *lsp-name* secondary *segment-list-name*]**.

NOTE: Starting in Junos OS Release 19.1R1, the **bfd-liveness-detection** statement is not supported at the **[edit protocols source-packet-routing [segment-list](#)]** hierarchy level.

Description

Bidirectional forwarding detection options.

Options

minimum-interval—Minimum transmit and receive interval (milliseconds).

Range: 1 through 255000

multiplier—Detection time multiplier.

Default: 3

Range: 1 through 255

no-router-alert-option—Do not set the Router Alert option in IP header.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing

class-of-service (Protocols MPLS)

Syntax

```
class-of-service class-of-service cos-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Class-of-service (CoS) value given to all packets in the LSP.

The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.

Options

cos-value—CoS value. A higher value typically corresponds to a higher level of service.

Range: 0 through 7

Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Class of Service for MPLS LSPs](#) | 1311

[Configuring the Ingress Router for Static LSPs | 604](#)[Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs | 607](#)

compute-options

Syntax

```
compute-options;
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Release Information

Statement introduced in Junos OS Release 19.2R1-S1 on MX Series routers.

Description

Configure compute options applicable to all the computed paths.

Required Privilege Level

routing

RELATED DOCUMENTATION

[Enabling Distributed CSPF for Segment Routing LSPs | 615](#)[compute-profile | 2575](#)

compute-profile

Syntax

```
compute-profile name {
  protected {
    mandatory;
  }
  unprotected {
    mandatory;
  }
  admin-group include-any [ include-any ... ]include-all [ include-all ... ]exclude [ exclude ... ];
  compute-segment-list compute-segment-list;
  maximum-computed-segment-lists maximum-computed-segment-lists;
  maximum-segment-list-depth maximum-segment-list-depth;
  metric-type {
    (igp | te);
  }
  no-label-stack-compression;
}
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Release Information

Statement introduced in Junos OS Release 19.2R1-S1 on MX Series Routers.

Description

Configure the compute profile for dynamically computed paths. You can use a compute profile to logically group the computation constraints. These compute profiles are referenced by the segment routing paths for computing the primary and secondary segment routing LSPs.

Options

name—Name of the computation-profile.

protected—Choose protected labels if available.

Values:

- **mandatory**—Mandatorily choose protected labels.

unprotected—Choose unprotected labels if available.

Values:

- **mandatory**—Mandatorily choose unprotected labels.

compute-segment-list—Name of the compute type segment list.

maximum-computed-segment-lists—Maximum number of segment-lists (ECMP paths) to be computed.

Range: 1 through 128

maximum-segment-list-depth—Maximum depth of computed path.

Range: 1 through 16

metric-type—Specify the metric type used for computation.

Values:

- **igp**—Interior gateway protocol metric.
- **te**—Traffic-engineering metric.

no-label-stack-compression—Provide fully expanded path, using adjacency segment identifiers.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing

RELATED DOCUMENTATION

[Enabling Distributed CSPF for Segment Routing LSPs](#) | 615

[compute-options](#) | 2574

connections (MPLS)

Syntax

```
connections {  
  remote-interface-switch connection-name {  
    interface interface-name.unit-number;  
    transmit-lsp label-switched-path;  
    receive-lsp label-switched-path;  
  }  
}
```

Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Define the connection between two circuits in a CCC connection.

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

[Example: Configuring MPLS on EX8200 and EX4500 Switches](#) | 59

constituent-list

Syntax

```
constituent-list constituent-list-name {
  (administrative-group [ group-names ] | administrative-group-extended [ extended-administrative-group-names ] |
  srlg [ srlg-names ]);
}
```

Hierarchy Level

```
[edit logical systems logical-systems-name protocols mpls]
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 17.1 for all platforms.

Description

Create a list of traffic engineering attributes called constituent attributes, which are the link and node attributes whose logical combination makes up an abstract hop. The constituent attributes are listed under administrative groups, extended administrative groups, and Shared Risk Link Groups (SRLGs).

Options

constituent-list-name—Name of the constituent list that includes constituent traffic engineering attributes for use in the abstract hop definition.

administrative-group [group-names]—Names of administrative groups to include in the constituent list.

administrative-group-extended [extended-administrative-group-names]—Names of extended administrative groups to include in the constituent list.

srlg [srlg-names]—Names of SRLGs to include in the constituent list.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring Abstract Hops for MPLS LSPs | 446](#)

[abstract-hop | 2534](#)

[show mpls abstract-hop-membership | 3170](#)

container-label-switched-path

Syntax

```
container-label-switched-path lsp-name {  
    disable;  
    description description;  
    label-switched-path-template;  
    splitting-merging;  
    suffix string;  
    to ip-address;  
}
```

Hierarchy Level

```
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

Description

Configure a multi-label-switched path (LSP) tunnel between the ingress and the egress routers. The container LSP consists of several member LSPs to the same destination.

Options

disable—Disable MPLS container-label-switched path.

description *description*—Text describing the container LSP.

suffix *string*—Suffix to generate names of member LSPs of the container LSP.

to *ip-address*—IP address of the egress router.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

corouted-bidirectional

Syntax

```
corouted-bidirectional;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify that the label-switched path be established as a corouted bidirectional packet LSP. You cannot configure this statement at the same time as the **corouted-bidirectional-passive** statement.

Default

This statement is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Corouted Bidirectional LSPs | 570](#)

[corouted-bidirectional-passive | 2581](#)

corouted-bidirectional-passive

Syntax

```
corouted-bidirectional-passive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify that the label-switched path be a passive LSP associated with a bidirectional LSP when it is signaled at the ingress router. This passive LSP enables the MPLS application to utilize the reverse LSP. You cannot configure this statement at the same time as the **corouted-bidirectional** statement.

Default

This statement is disabled by default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Corouted Bidirectional LSPs | 570](#)

[corouted-bidirectional | 2580](#)

credibility

Syntax

```
credibility {  
    direct;  
    isis-level-1;  
    isis-level-2;  
    ospf;  
    static;  
    unknown;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering database export],  
[edit protocols mpls traffic-engineering database export]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced in Junos OS Release 17.1R1 on QFX Series and QFX10000 switches.

Description

Configure preference values for entries from BGP-TE to the traffic engineering database. A protocol with a higher credibility value is preferred over a protocol with a lower credibility value.

The credibility order for the BGP-TE protocols is as follows:

- **Unknown**—80
- **OSPF**—81
- **ISIS Level 1**—82
- **ISIS Level 2**—83
- **Static**—84
- **Direct**—85

Options

direct— Entries sourced from directly connected links.

isis-level-1—Entries sourced from IS-IS Level 1.

isis-level-2—Entries sourced from IS-IS Level 2.

ospf—Entries sourced from OSPF.

static—Entries sourced from static configuration.

unknown—Entries sourced from unknown entities.

Range: 0 through 512

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [traffic-engineering](#) | **2818**

database

Syntax

```
database {  
    export;  
    import;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering],  
[edit protocols mpls traffic-engineering]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced in Junos OS Release 17.1R1 for QFX Series and QFX10000 switches.

Description

Include link and node entries from the traffic engineering database into the **lsdist.0** routing information base (RIB), so it gets picked up by the BGP export policy.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [traffic-engineering](#) | **2818**

delay (querier)

Syntax

```
delay {
  traffic-class tc-value {
    average-sample-size sample size;
    padding-size size;
    query-interval milliseconds;
    rtt-delay-threshold rtt threshold value;
    twcd-delay-threshold twcd threshold value;
  }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure delay measurement options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements | 273](#)

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

delay (responder)

Syntax

```
delay {
  min-query-interval milliseconds;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring responder]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure delay measurement options.

Options

min-query-interval *milliseconds*—(Optional) Specify the minimum query interval that the responder supports. If the minimum query interval of the responder is greater than the query interval configured at querier, the effective message query rate will be the minimum query interval configured for the responder.

Default: 10 seconds

Range: 1000 through 4294967295 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements](#) | 273

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview](#) | 243

[performance-monitoring \(Protocols MPLS\)](#) | 2727

description (Protocols MPLS)

Syntax

```
description text;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Provides a textual description of the LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show mpls lsp detail** command and has no effect on the operation of the LSP.

Options

text—Provide a textual description of the LSP. The description text can be no more than 80 characters in length.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring a Text Description for LSPs](#) | 528

description (Protocols Layer 2 VPN)

Syntax

```
description text;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn site site-name
  interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Describe the VPN or virtual private LAN service (VPLS) routing instance.

Options

text—Provide a text description. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show route instance detail** command and has no effect on operation.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Site

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

deselect-on-bandwidth-failure

Syntax

```
deselect-on-bandwidth-failure {  
    tear-lsp;  
}
```

Hierarchy Level

```
[edit protocols mpls],  
[edit protocols mpls label-switched-path path-name]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Deselect an active path if it does not meet the auto-bandwidth criteria required for path selection. The **deselect-on-bandwidth-failure** statement does not apply to static bandwidth.

Options

tear-lsp— Bring down an active path if none of the paths are able to reserve the required bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

diffserv-te

Syntax

```
diffserv-te {
  bandwidth-model {
    extended-mam;
    mam;
    rdm;
  }
  te-class-matrix {
    tnumber {
      priority priority;
      traffic-class {
        cnumber priority priority;
      }
    }
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify properties for differentiated services in traffic engineering.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Routers for DiffServ-Aware Traffic Engineering](#) | 1199

disable (Protocols MPLS)

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit logical-systems logical-system-name protocols mpls interface interface-name],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls],  
[edit protocols mpls interface interface-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Disable the functionality of the configured object.

Default

The configured object is enabled (operational) unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[label-switched-path](#) | [2653](#)

[interface](#) | [2648](#)

dual-transport

Syntax

```
dual-transport {  
    inet-lsr-id inet-lsr-id;  
    inet6-lsr-id inet6-lsr-id;  
}
```

Hierarchy Level

[edit protocols ldp]

Release Information

Statement introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

Description

Configure to allow Junos LDP to establish the TCP connection over IPv4 with IPv4 neighbors, and over IPv6 with IPv6 neighbors as a single-stack LSR. **inet-lsr-id** and **inet6-lsr-id** are the two LSR IDs that have to be configured to establish an LDP session over IPv4 and IPv6 TCP transport. These two IDs should be non-zero and must be configured with different values.

Options

inet-lsr-id *inet-lsr-id*— Configure the LSR ID to establish an LDP session over IPv4 TCP transport.

inet6-lsr-id *inet6-lsr-id*— Configure the LSR ID to establish an LDP session over IPv6 TCP transport.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[LDP Native IPv6 Support Overview | 932](#)

[Example: Configuring LDP Native IPv6 Support | 1054](#)

[Configuring LDP Native IPv6 Support | 1053](#)

[family \(Protocols LDP\) | 2614](#)

dynamic (Source Packet Routing)

Syntax

```
dynamic {
  protected {
    mandatory;
  }
  unprotected {
    mandatory;
  }
}
```

Hierarchy Level

[edit protocols source-packet-routing [segment-list](#)]

Release Information

Statement introduced in Junos OS Release 19.2R1 on all platforms.

Description

Enable dynamic computation of segment routing label-switched paths (LSPs) based on tunnel destination and translation service to fetch the corresponding segment identifiers (SIDs).

NOTE:

When the dynamic statement is enabled, all the next hops must have an IP address assigned as a minimum configuration. In the case of segment-lists, if a next hop has both IP address and label configured, then the configured label is retained.

NOTE:

- Because translation service use IGP instance of traffic-engineered database (TED), you must include the **igp-topology** statement at the **[edit protocols isis traffic-engineering]** hierarchy level for successful translation.
- The **auto-translation** and **dynamic** statements are mutually exclusive, and only either of them can be configured under a segment-list.

Options

protected—(Optional) Ensures a protected adjacency SID is used for the links of the LSP.

unprotected—(Optional) Ensures an unprotected adjacency SID is used for the links of the LSP.

mandatory—(Optional) Enabled translation to fail if the specified type of adjacency SID cannot be found for a link. This option does not have effect on node SIDs.

Required Privilege Level

routing

RELATED DOCUMENTATION

| [Understanding Static Segment Routing LSP in MPLS Networks](#) | 632

dynamic-tunnels

Syntax

```
dynamic-tunnels tunnel-name {
  destination-networks prefix;
  gre;
  rsvp-te entry-name {
    destination-networks network-prefix;
    label-switched-path-template (Multicast) {
      default-template;
      template-name;
    }
  }
  source-address address;
  spring-te;
  traceoptions;
  tunnel-attributes name {
    dynamic-tunnel-anchor-pfe dynamic-tunnel-anchor-pfe;
    dynamic-tunnel-anti-spoof (off | on);
    dynamic-tunnel-mtu dynamic-tunnel-mtu;
    dynamic-tunnel-source-prefix dynamic-tunnel-source-prefix;
    dynamic-tunnel-type V4oV6;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options],
[edit logical-systems logical-system-name routing-options],
[edit routing-instances routing-instance-name routing-options],
[edit routing-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure a dynamic tunnel between two PE routers.

NOTE: ACX Series routers do not support the **gre** statement.

Configure dynamic IPv4-over-IPv6 tunnels and define their attributes to forward IPv4 traffic over an IPv6-only network. IPv4 traffic is tunneled from customer premises equipment to IPv4-over-IPv6 gateways. You must also configure **extended-nexthop** option at **[edit protocols bgp family inet unicast]** hierarchy level to allow BGP to route IPv4 address families over an IPv6 session.

Options

gre—Enable dynamic generic routing encapsulation type tunnel mode for IPv4

Values:

- **next-hop-based-tunnel**—Enable next hop base dynamic-tunnel for steering IPv4 traffic with IPv6 next hop address.

source-address—Specify the source address of the tunnel.

tunnel-name—Name of the dynamic tunnel.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>extended-nexthop</i>
<i>tunnel-attributes</i>
<i>Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks</i>
<i>Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP</i>

egress-protection (MPLS)

Syntax

```
egress-protection {
  context-identifier context-id {
    primary | protector;
    metric igp-metric-value;
    advertise-mode (stub-alias | stub-proxy);
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Options **primary**, **protector**, and **metric** introduced in Junos OS Release 11.4R3.

Option **advertise-mode** introduced in Junos OS Release 13.3.

Description

Enables an Edge Protection Virtual Circuit (EPVC) for the MPLS protocol.

Options

context-identifier *context-id-ip-address*—(Optional) The context identifier IPv4 address.

metric *igp-metric-value*—(Optional) The IGP metric value ranging from **2** through **16777215**.

(primary | protector)—On the primary PE router, configure as type **primary**. On the protector PE router, configure as type **protector**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Egress Protection for Layer 3 VPN Services

| *Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP*

encapsulation-type (Layer 2 VPNs)

Syntax

```
encapsulation-type (atm-aal5 | atm-cell | atm-cell-port-mode | atm-cell-vc-mode | atm-cell-vp-mode | cesop |
  cisco-hdlc | ethernet | ethernet-vlan | frame-relay | frame-relay-port-mode | interworking | ppp | satop-e1 | satop-e3
  | satop-t1 | satop-t3);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols l2circuit neighbor address interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls],
[edit protocols l2circuit neighbor address interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn neighbor address],
[edit routing-instances routing-instance-name protocols vpls],
[edit routing-instances routing-instance-name protocols vpls neighbor address]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Specify the type of Layer 2 traffic originating from the CE device. Only the **ethernet** and **ethernet-vlan** encapsulation types are supported for VPLS. Not all encapsulation types are supported on the switches. See the switch CLI.

Options

atm-aal5—ATM Adaptation Layer (AAL/5)

atm-cell—ATM cell relay

atm-cell-port-mode—ATM cell relay port promiscuous mode

atm-cell-vc-mode—ATM VC cell relay nonpromiscuous mode

atm-cell-vp-mode—ATM virtual path (VP) cell relay promiscuous mode

cesop—CESOP-based Layer 2 VPN

cisco-hdlc—Cisco Systems-compatible HDLC

ethernet—Ethernet

ethernet-vlan—Ethernet VLAN

frame-relay—Frame Relay

frame-relay-port-mode—Frame Relay port mode

interworking—Layer 2.5 interworking VPN

ppp—PPP

satsop-e1—SATSOP-E1-based Layer 2 VPN

satsop-e3—SATSOP-E3-based Layer 2 VPN

satsop-t1—SATSOP-T1-based Layer 2 VPN

satsop-t3—SATSOP-T3-based Layer 2 VPN

Default: For VPLS networks, the default encapsulation type is **ethernet**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Encapsulation Type

Configuring VPLS Routing Instances

Configuring the Encapsulation Type for the Layer 2 Circuit Neighbor Interface

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

encoding-type

Syntax

```
encoding-type (ethernet | packet | pdh | sonet-sdh);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-attributes],  
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the encoding type of payload carried by the LSP. It can be any of the following:

- **ethernet**—Ethernet
- **packet**—Packet
- **pdh**—Plesiochronous digital hierarchy (PDH)
- **sonet-sdh**—SONET/SDH

Default

packet

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Encoding Type](#) | 1373

entropy-label

Syntax

```
entropy-label {  
    ingress-policy ingress-policy-name;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
```

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Assists the transit router in load-balancing MPLS traffic across ECMP paths or Link Aggregation groups by introducing the entropy label to the MPLS label stack. The entropy label allows routers to load balance MPLS traffic by using a hash-input without the need to perform deep packet inspection. Deep packet inspection requires more of the router's processing power and is not a capability shared by all routers.

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Entropy Label for LSPs](#) | 573

entropy-label

Syntax

```
entropy-label {  
    import policy-name;  
    no-next-hop-validation;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system name protocols bgp family inet labeled-unicast],  
[edit logical-systems logical-system name protocols bgp group group-name family inet labeled-unicast],  
[edit logical-systems logical-system name protocols bgp group group-name neighbor address labeled-unicast],  
[edit protocols bgp family inet labeled-unicast],  
[edit protocols bgp group group-name family inet labeled-unicast],  
[edit protocols bgp group group-name neighbor address labeled-unicast]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Insert the entropy label into the BGP labeled unicast LSP packets, which assists the transit router in load-balancing BGP traffic across equal-cost multipaths or link aggregation groups. The ingress label edge router inspects the flow information of a packet and maps it to an entropy label, which is inserted into the BGP label stack. LSRs in the core use this entropy label as the key to hash the packet and direct it to the correct path.

Options

import *policy-name*— (Optional) Specify a policy that lists the routes that allow the use of entropy labels.

no-next-hop-validation—(Optional) Do not validate the next-hop field in the entropy label capability attribute against the route next hop.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

labeled-unicast

[policy-statement](#) | [2733](#)

Configuring an Entropy Label for a BGP Labeled Unicast LSP

[Example: Configuring an Entropy Label for a BGP Labeled Unicast LSP](#) | [575](#)

Understanding Entropy Label for BGP Labeled Unicast LSP

ethernet-vlan (Protocols Link Management)

Syntax

```
ethernet-label {  
    vlan-id-range vlan-id-range;  
}
```

Hierarchy Level

```
[edit protocols link-management te-link te-link-name]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Specify the TE-link to be used for Layer 2 VLAN label-switched path (LSP).

Options

vlan-id-range *vlan-id-range*—Pool of VLAN IDs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

ether-pseudowire

Syntax

```
ether-pseudowire {  
    zero-control-word;  
}
```

Hierarchy Level

[edit forwarding-options enhanced-hash-key family mpls]

Release Information

Statement introduced in Junos OS Release 16.1 for the MX Series.

Description

Load-balance IP over Ethernet pseudowire. Presence of zero-control-word in the payload indicates an Ethernet frame.

zero-control-word— Precedes Ethernet packet to indicate the start Ethernet frame in an MPLS ether-pseudowire payload.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[enhanced-hash-key](#)

[hash-key](#)

[family mpls | 2615](#)

[MPLS Encapsulated Payload Load-balancing Overview | 239](#)

exclude (for Administrative Groups)

Syntax

```
exclude [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name admin-group],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary | secondary) path-name  
  admin-group],  
[edit protocols mpls label-switched-path lsp-name admin-group],  
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name admin-group]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define the administrative groups to exclude for an LSP or for a path's primary and secondary paths.

Options

group-names—Names of one or more groups defined with the **admin-groups** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Administrative Groups for LSPs](#) | 532

exclude (for Fast Reroute)

Syntax

```
(exclude [ group-names ] | no-exclude);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],  
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Control exclusion of administrative groups:

- **exclude**—Define the administrative groups to exclude for fast reroute.
- **no-exclude**—Disable administrative group exclusion.

Options

group-names—Names of one or more groups defined with the **admin-groups** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute | 479](#)

[admin-groups | 2548](#)

exclude-srlg

Syntax

```
exclude-srlg;
```

Hierarchy Level

```
[edit protocols mpls],  
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls label-switched-path path-name],  
[edit logical-systems logical-system-name protocols mpls label-switched-path path-name],  
[edit protocols rsvp interface interface-name link-protection],  
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],  
[edit protocols rsvp interface interface-name link-protection bypass destination],  
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass destination]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

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.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Excluding SRLG Links Completely for the Secondary LSP](#) | 290

exp

Syntax

```
exp classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-points [aliases] [3-bit-patterns];
    }
  }
}
```

Hierarchy Level

```
[edit class-of-service classifiers],
  [edit class-of-service code-point-aliases],
  [edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules],
  [edit class-of-service rewrite-rules]
```

Release Information

Statement introduced in Junos OS Release 10.1 for EX Series switches.

Description

Define the experimental bits (EXP) code point mapping that is applied to MPLS packets. You can define an **exp** classifier only on EX3200 switches, EX4200 and EX8200 standalone switches, and EX8200 Virtual Chassis. You can bind an **exp** rewrite rule on EX8200 standalone switches and EX8200 Virtual Chassis.

EX Series switches support only one EXP code mapping on the switch (either default or custom). It is applied globally and implicitly to all the MPLS-enabled interfaces on the switch. You cannot bind it or disable it on individual interfaces.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Using CoS with MPLS Networks on EX Series Switches | 1324](#)

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

[Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) | 1323](#)

expand-loose-hop

Syntax

```
expand-loose-hop;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 7.6.

Point-to-multipoint LSP support introduced in Junos OS Release 11.2.

Description

Allow an LSP to traverse multiple OSPF areas within a service provider's network.

Allows a point-to-multipoint LSP to span multiple domains in a network. Effectively, this allows you to configure one or more sub-LSPs (branches) in separate network domains. Examples of such domains include OSPF areas and autonomous systems (ASs). A sub-LSP of an inter-domain point-to-multipoint LSP can be intra-area, inter-area, or inter-AS, depending on the location of the egress node (leaf) with respect to the ingress node (source). Only OSPF areas are supported for inter-domain point-to-multipoint LSPs. IS-IS levels are not supported.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Enabling Interarea Traffic Engineering | 1140](#)

[Configuring Inter-Domain Point-to-Multipoint LSPs | 722](#)

explicit-null (Protocols MPLS)

Syntax

```
explicit-null;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Advertise label 0 to the egress router of an LSP.

Default

If you do not include the **explicit-null** statement in the MPLS configuration, label 3 (implicit null) is advertised.

NOTE: Junos OS does not support explicit null routes with next hops to virtual tunnel (vt-) interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring RSVP to Pop the Label on the Ultimate-Hop Router](#) | 858

export (MPLS Traffic engineering database)

Syntax

```
export {  
  credibility;  
  policy policy-name;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering database],  
[edit protocols mpls traffic-engineering database]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Configure the traffic engineering database export-related parameters.

Options

policy *policy-name*—Name of the export policy.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [traffic-engineering](#) | **2818**

failure-action (Protocols MPLS)

Syntax

```
failure-action {  
    make-before-break teardown-timeout seconds;  
    teardown;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls oam bfd-liveness-detection],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name oam bfd-liveness-detection],  
[edit protocols mpls label-switched-path lsp-name oam bfd-liveness-detection],  
[edit protocols mpls oam bfd-liveness-detection]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Configure route and next-hop properties in the event of a Bidirectional Forwarding Detection (BFD) protocol session failure event on an RSVP label-switched path (LSP). The failure event could be an existing BFD session that has gone down or a BFD session that never came up. RSVP adds back the route or next hop when the relevant BFD session comes back up.

Options

make-before-break—When a BFD session fails for an RSVP LSP, an attempt is made to signal a new LSP path before tearing down the old LSP path.

teardown—When a BFD session fails for an RSVP LSP, the associated LSP path is taken down and resigaled immediately.

teardown-timeout *seconds*—When you configure the **make-before-break** option, you can specify a time in seconds for the **teardown-timeout** option. At the end of the time specified, the associated RSVP LSP is automatically torn down and resigaled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

family

Syntax

```
family {  
    inet;  
    inet6;  
}
```

Hierarchy Level

```
[edit protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

Description

Configure the address family as inet for IPv4 or inet6 for IPv6, or both. If the address family is not configured, then the default address family is IPv4.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[LDP Native IPv6 Support Overview | 932](#)

[Example: Configuring LDP Native IPv6 Support | 1054](#)

[Configuring LDP Native IPv6 Support | 1053](#)

[dual-transport | 2592](#)

family mpls

Syntax

```
family mpls {
  all-labels;
  label-1;
  label-2;
  label-3;
  no-labels;
  no-label-1-exp;
  payload {
    ether-pseudowire {
      zero-control-word;
    }
    ip {
      disable;
      layer-3-only;
      port-data {
        source-msb;
        source-lsb;
        destination-msb;
        destination-lsb;
      }
    }
  }
}
```

Hierarchy Level

[edit forwarding-options hash-key]

Release Information

Statement introduced before Junos OS Release 7.4.

no-label-1-exp option introduced in Junos OS Release 8.0.

label-3 and **no-labels** options introduced in Junos OS Release 8.1.

ether-pseudowire statement introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.

all-labels and **payload ip disable** options introduced in Junos OS Release 12.1X48R2. (PTX Series Packet Transport Routers only).

zero-control-word option introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

Description

For aggregated Ethernet and SONET/SDH interfaces only, configure load balancing based on MPLS labels and payload. Only the IPv4 protocol is supported.

Options

family mpls—(Aggregated Ethernet interfaces, aggregated SONET/SDH interfaces, and multiple equal-cost MPLS next hops only) Incorporate MPLS label and payload information into the hash key for per-flow load balancing. Only the IPv4 protocol is supported.

- **all-labels**—(PTX Series Packet Transport Routers only) Up to eight MPLS labels are included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. This is the default setting.
- **label-1**—(M120, M320, MX Series, and T Series routers only) Include the first MPLS label into the hash key. This is used for a one-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels.
- **label-2**—(M120, M320, MX Series, and T Series routers only) Include the second MPLS label into the hash key. This is used for a two-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels. To use the second MPLS label in the hash key, include both the **label-1** and **label-2** statements at the **[edit forwarding-options hash-key family mpls]** hierarchy level. By default, the router provides hashing on the first and second labels. If both labels are specified, the entire first label and the first 16 bits of the second label are hashed.
- **label-3**—(M120, M320, MX Series, and T Series routers only) Include the third MPLS label into the hash key. To use the third MPLS label, include the **label-1**, **label-2**, and **label-3** statements at the **[edit forwarding-options hash-key family mpls]** hierarchy level.
- **no-labels**—Include no MPLS labels into the hash key.
- **no-label-1-exp**—(M120, M320, MX Series, and T Series routers only) The EXP bit of the first label is not used in the hash calculation to avoid reordering complications.
- **payload**—Incorporate bits from the IP payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
 - **disable**—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
 - **ether-pseudowire**—(M120, M320, MX Series, and T Series routers only) Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
 - **zero-control-word**—(M Series, MX Series, and PTX Series) Precedes Ethernet packet to indicate the start of an Ethernet frame in an MPLS ether-pseudowire payload.
 - **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels. For the PTX Series Packet Transport Routers, this is the default setting with both Layer 3 and Layer 4 IP information included in the hash key.
 - **disable**—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
 - **layer-3-only**—Include only Layer 3 IP information from the IP payload data into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
 - **port-data**—(M120, M320, MX Series, and T Series routers only) Include the source and destination port field information into the hash key. By default, the most significant byte and least significant byte of the source and destination port fields are hashed. To select specific bytes to be hashed,

include one or more of the **source-msb**, **source-lsb**, **destination-msb**, and **destination-lsb** options at the **[edit forwarding-options hash-key family mpls payload ip port-data]** hierarchy level. To prevent all four bytes from being hashed, include the **layer-3-only** statement at the **[edit forwarding-options hash-key family mpls payload ip]** hierarchy level.

- **destination-lsb**—Include the least-significant byte of the destination port.
- **destination-msb**—Include the most-significant byte of the destination port.
- **source-lsb**—Include the least-significant byte of the source port.
- **source-msb**—Include the most-significant byte of the source port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Load Balancing Based on MPLS Labels | 490](#)

[Configuring Load Balancing for Ethernet Pseudowires | 1308](#)

fast-reroute (Protocols MPLS)

Syntax

```
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);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Establish detours for the LSP so that if a node or link in the LSP fails, the traffic on the LSP can be rerouted with minimal packet loss.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute](#) | 479

[Fast Reroute Overview](#) | 477

[MPLS Feature Support on QFX Series and EX4600 Switches](#) | 26

[Understanding Interprovider and Carrier-of-Carriers VPNs](#) | 1647

fate-sharing

Syntax

```
fate-sharing {
  group group-name {
    cost value;
    from address <to address>;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options],
[edit routing-options],
[edit routing-instances routing-instance-name routing-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify a backup path in case the primary path becomes unusable.

You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.

Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.

Options

cost *value*—Cost assigned to the group.

Range: 1 through 65,535

Default: 1

from *address*—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate **from** addresses in the group.

group *group-name*—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.

to *address*—(Optional) Address of egress router. For point-to-point link objects, you must specify both a **from** and a **to** address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Alternate Backup Paths Using Fate Sharing | 518](#)

MPLS Applications User Guide

forwarding-rib

Syntax

```
forwarding-rib name {  
    inet-import [ inet-import ... ];  
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options dynamic-tunnels],  
[edit logical-systems name routing-options dynamic-tunnels],  
[edit logical-systems name tenants name routing-instances name routing-options dynamic-tunnels],  
[edit routing-instances name routing-options dynamic-tunnels],  
[edit routing-options dynamic-tunnels],  
[edit tenants name routing-instances name routing-options dynamic-tunnels]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on PTX Series routers and QFX Series switches.

Description

Configure policy control for forwarding routing table next hops for MPLS-over-UDP dynamic tunnels. With this configuration, you can resolve the dynamic tunnel destination routes over select prefixes.

Options

name—Name of the routing table.

inet-import—Name of the import policy for IPv4 dynamic-tunnels.

Required Privilege Level

routing

RELATED DOCUMENTATION

| [Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels](#) | 367

forwarding-table

Syntax

```
forwarding-table {  
  export [ policy--names ];  
  (indirect-next-hop | no-indirect-next-hop);  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],  
[edit routing-options]
```

Release Information

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure information about the routing device's forwarding table.

The remaining statements are explained separately. See [CLI Explorer](#).

Options

remnant-holdtime—Sets the remnant hold time, which is required for the MXVC-ISSU, where the recommended value is 900..

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring Per-Packet Load Balancing*

from (Protocols MPLS)

Syntax

```
from address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Specify the source address to use for the LSP.

The address you specify does not affect the outgoing interface used by the LSP.

Default

If you do not include this statement, the software automatically selects the loopback interface as the address.

Options

address—IP address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Ingress Router Address for LSPs](#) | 514

gpId

Syntax

```
gpId { ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 | pos-scrambling-crc-32 |
      pos-no-scrambling-crc-32 | ppp};
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-attributes],
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Release Information

Statement introduced before Junos OS Release 7.4.

pos-scrambling-crc-16, **pos-no-scrambling-crc-16**, **pos-scrambling-crc-32**, and **pos-no-scrambling-crc-32** options added in Junos OS Release 8.0.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the type of payload carried by the LSP. It can be any of the following:

- **ethernet**—Ethernet (GPID value: 33)
- **hdlc**—High-level Data Link Control (HDLC) (GPID value: 44)
- **ipv4**—IP version 4 (GPID value: 0x0800)
- **pos-no-scrambling-crc-16**—for interoperability with other vendors' equipment (GPID value: 29)
- **pos-no-scrambling-crc-32**—for interoperability with other vendors' equipment (GPID value: 30)
- **pos-scrambling-crc-16**—for interoperability with other vendors' equipment (GPID value: 31)
- **pos-scrambling-crc-32**—for interoperability with other vendors' equipment (GPID value: 32)
- **ppp**—Point-to-Point Protocol (PPP) (GPID value: 50)

Default

ipv4

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

gre (Routing Options)

Syntax

```
gre;  
next-hop-based-tunnel;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],  
[edit routing-options dynamic-tunnels tunnel-name]
```

Release Information

Statement introduced in Junos OS Release 10.1.

next-hop-based-tunnel option introduced in Junos OS Release 16.2.

Description

Enable generic routing encapsulation (GRE) type for IPv4 to automatically establish label-switched paths (LSPs) for any new provider edge (PE) router added to a full mesh of LSPs.

Options

next-hop-based-tunnel—Create a tunnel composite next hop for every dynamic GRE tunnel configured.

The tunnel composite next hop includes the dynamic tunnel's encapsulation data and a VPN label (when chained composite next hop is not enabled). When this option is not configured, the default interface-based tunnel mode is enabled. By configuring this option, a device can scale up to 32,000 IP tunnels, which is otherwise restricted to the system limit on the number of interfaces supported.

At a given point in time, either the next-hop-based dynamic tunnel or the default interface-based dynamic GRE tunnel can exist on a device. A switchover from one tunnel mode to another deletes the existing tunnels and creates new tunnels in the new tunnel mode. As a result, a tunnel mode switchover can impact network performance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS-Signaled LSPs to Use GRE Tunnels](#) | 131

hop-limit

Syntax

```
hop-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name fast-reroute],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:

- LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.
- Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.
- Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.

Options

number—Maximum number of hops.

Range: 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)

Default: 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute | 479](#)

[Limiting the Number of Hops in LSPs | 546](#)

[Configuring the Hop Limit for Bypass LSPs | 875](#)

import (MPLS Traffic Engineering Database)

Syntax

```
import {
  bgp-ls-identifier domain-identifier;
  identifier identifier;
  policy policy-name;
  igp-topology{
    bgp-link-state;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls traffic-engineering database],
[edit protocols mpls traffic-engineering database]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Configure the traffic engineering database import parameters.

Options

bgp-ls-identifier *domain-identifier*—BGP-TE domain identifier.

identifier *identifier*—BGP-TE identifier.

Range: 2 through 18446744073709551615

policy *policy-name*—Name of the import policy.

igp-topology—Download IGP topology information into the traffic engineering database (TED). In Junos OS, the IGP installs topology information into a database called the traffic engineering database. The traffic engineering database contains the aggregated topology information. The IGP routes are installed by the traffic engineering database on behalf of the corresponding IGP into a user-visible routing table called `Isdist.0`, subject to route policies.

bgp-link-state—Export IGP topology information into BGP-Link State (BGP-LS) from the `Isdist.0` routing table. The `Isdist.0` routing table stores the network topology information from the traffic engineering database. The BGP-LS reads IGP entries from `Isdist.0` and advertises the information to BGP peers.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [traffic-engineering](#) | **2818**

ip-tunnel-rpf-check

Syntax

```
ip-tunnel-rpf-check {
    mode (strict | loose);
    fail-filter filter-name;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name routing-options forwarding-table]
```

Release Information

Statement introduced in Junos OS Release 17.1 for MX Series Routers with MICs.

Description

Configure the system to enable anti-spoofing protection for next-hop-based dynamic tunnels, where reverse path forwarding checks are placed to ensure that the tunnel traffic is received from a legitimate source through designated IP tunnel, where the source is reachable on the same tunnel on which the packet was received.

When a packet comes from a nondesignated source, the reverse path forwarding check fails in the strict mode, and passes in the loose mode. When a packet comes from a nonexistent source, the reverse path forwarding check fails.

By default, the reverse path forwarding check is in strict mode, where the packets are not forwarded if the source of the packet is from a nondesignated tunnel.

Options

mode (strict | loose)—(Optional) Specify the mode of the reverse path forwarding check to enable anti-spoofing protection for next-hop-based dynamic tunnels.

In the strict mode (default), the reverse path forwarding check fails when the packet is received from a nondesignated tunnel source. The check passes only for packets from designated tunnels.

In the loose mode, the reverse path forwarding check passes even if the packet is received from a nondesignated tunnel source.

When the packet is from a nonexistent tunnel source, the reverse path forwarding check fails in both the strict and loose modes.

Default: If you omit the mode statement, the default behavior is strict mode.

fail-filter *filter-name*—(Optional) Attach a filter to the Layer 3 VPN to log packets that failed the reverse path forwarding check.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview | 385](#)

[Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels | 388](#)

ipv4 (Family MPLS)

Syntax

```

ipv4 {
  destination-address ip-address {
    except;
  }
  destination-prefix-list destination-prefix-list-name {
    except;
  }
  protocol protocol {
    (source-port | source-port-except);
    (destination-port | destination-port-except);
  }
  source-address ip-address {
    except;
  }
  source-prefix-list source-prefix-list-name {
    except;
  }
}

```

Hierarchy Level

```
[edit firewall family mpls filter name term name from ip-version]
```

Release Information

Statement introduced in Junos OS Release 18.4R1 on MX Series routers with MPC and MIC interfaces.

Description

Define Layer 3 and Layer 4 match items to match IPv4 packets for IP-based filtering of MPLS traffic.

Options

destination-address *ip-address*—Match MPLS traffic with the specified IPv4 destination address.

destination-prefix-list *destination-prefix-list-name*—Match MPLS traffic with the specified IPv4 destination prefixes. The prefix-list is defined under the **[edit policy-options prefix-list *prefix-list-name*]** hierarchy level.

protocol *protocol*—Specify one or a range of inner IPv4 protocols for IP-based filtering of MPLS traffic.

source-address *ip-address*—Match MPLS traffic with the specified IPv4 source address.

source-prefix-list *source-prefix-list-name*—Match MPLS traffic with the specified IPv4 source prefixes. The prefix-list is defined under the **[edit policy-options prefix-list *prefix-list-name*]** hierarchy level.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic](#) | 135

ipv6 (Family MPLS)

Syntax

```

ipv6 {
  destination-address destination-ip-address {
    except;
  }
  destination-prefix-list prefix-list-name {
    except;
  }
  protocol protocol {
    (source-port | source-port-except);
    (destination-port | destination-port-except);
  }
  source-address ip-address {
    except;
  }
  source-prefix-list source-prefix-list-name {
    except;
  }
}

```

Hierarchy Level

```
[edit firewall family mpls filter name term name from ip-version],
```

Release Information

Statement introduced in Junos OS Release 18.4R1 on MX Series routers with MPC and MIC interfaces.

Description

Define Layer 3 and Layer 4 match items to match IPv6 packets for IP-based filtering of MPLS traffic.

Options

destination-address *ip-address*—Match MPLS traffic with the specified IPv6 destination address.

destination-prefix-list *destination-prefix-list-name*—Match MPLS traffic with the specified list of IPv6 destination prefixes. The prefix-list is defined under the **[edit policy-options prefix-list *prefix-list-name*]** hierarchy level. You must configure separate prefix-lists for IPv4 and IPv6 addresses.

protocol *protocol*—Specify one or a range of inner IPv6 next header for IP-based filtering of MPLS traffic.

source-address *ip-address*—Match MPLS traffic with the specified IPv6 source address.

source-prefix-list *source-prefix-list-name*—Match MPLS traffic with the specified IPv6 source prefixes. The prefix-list is defined under the **[edit policy-options prefix-list *prefix-list-name*]** hierarchy level. You must configure separate prefix-lists for IPv4 and IPv6 addresses.

Required Privilege Level

firewall

ip-version (Family MPLS)

Syntax

```
ip-version {
  ipv4;
  ipv6;
}
```

Hierarchy Level

```
[edit firewall family mpls filter name term name from]
```

Release Information

Statement introduced in Junos OS Release 18.4R1 on MX Series routers with MPC and MIC interfaces.

Description

Specify inner IP version to enable IP-based filtering of MPLS family filter.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic](#) | 135

include-all (for Administrative Groups)

Syntax

```
include-all [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name admin-group],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary | secondary) path-name  
  admin-group],  
[edit protocols mpls label-switched-path lsp-name admin-group],  
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name admin-group]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Require the LSP to traverse links that include all of the defined administrative groups.

Options

group-names—One or more names of groups defined with the **admin-groups** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Administrative Groups for LSPs | 532](#)

[admin-groups | 2548](#)

include-all (for Fast Reroute)

Syntax

```
(include-all [ group-names ] | no-include-all);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],  
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Control inclusion of administrative groups:

- **include-all**—Define the administrative groups that must all be included for fast reroute.
- **no-include-all**—Disable administrative group inclusion.

Options

group-names—One or more names of groups defined with the **admin-groups** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute](#) | 479

include-any (for Administrative Groups)

Syntax

```
include-any [ group-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name admin-group],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary | secondary) path-name  
  admin-group],  
[edit protocols mpls label-switched-path lsp-name admin-group],  
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name admin-group]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define the administrative groups to include for an LSP or for a path's primary and secondary paths.

Options

group-names—One or more names of groups defined with the **admin-groups** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Administrative Groups for LSPs](#) | 532

include-any (for Fast Reroute)

Syntax

```
(include-any [ group-names ] | no-include-any);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],  
[edit protocols mpls label-switched-path lsp-name fast-reroute]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Control inclusion of administrative groups:

- **include-any**—Define the administrative groups to include for fast reroute.
- **no-include-any**—Disable administrative group inclusion.

Options

group-names—One or more names of groups defined with the **admin-groups** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute](#) | 479

ingress (LSP)

Syntax

```
ingress {
  bandwidth bps;
  class-of-service cos-value;
  description string;
  entropy-label;
  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;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 10.1.

entropy-label option introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Configure an ingress LSR for a static LSP.

The remaining statements are explained separately

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Static LSPs](#) | 603

install (Protocols MPLS)

Syntax

```
install {
    destination-prefix <active>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name ingress]
[edit protocols source-packet-routing source-routing-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Support at the **[edit protocols source-packet-routing source-routing-path *lsp-name*]** hierarchy level introduced in Junos OS Release 19.1R1.

Description

Associate one or more prefixes with an LSP. When the LSP is up, all the prefixes are installed as entries into the inet.3 or inet6.3 routing table.

The support of the **install** statement at the **[edit protocols source-packet-routing source-routing-path *lsp-name*]** is applicable for both colored and non-colored segment routing LSPs.

For static colored LSPs, when the install prefix is configured, a route similar to the tunnel ingress route is installed in the **inetcolor.0** or **inet6color.0** routing table.

On the other hand, for static non-colored LSPs, when the install prefix is configured, a route similar to the tunnel **to** route is installed in **inet.3** or **inet6.3** routing table.

You can use the **show route table** command to view the **install** routes for both colored and non-colored segment routing traffic-engineered LSPs.

NOTE:

Take the following into consideration when configuring the **install *destination-prefix*** statement at the **[edit protocols source-packet-routing source-routing-path *lsp-name*]** hierarchy level:

- The install prefixes should be unique in a tunnel, and should not be the same as the tunnel **to** address. A commit check is done to ensure that the prefixes are unique.
- If two install prefixes are same across two different tunnels, then the gateways of the both tunnels are considered only if the segment lists are of the same resolution type. If the first hop resolution types vary, the route is not installed. In such cases, a system log message is generated to record the error.

For example:

```
[edit protocols source-packet-routing]
segment-list path-1 {
  hop-1 ip-address 172.0.12.2;
  hop-2 label 1000012;
  hop-3 label 1000013;
  hop-4 label 1000014;
}
source-routing-path lsp1 {
  to 10.10.10.1;
  install 20.20.20.2;
  install 30.30.30.3;
  primary {
    path-1;
  }
}
```

The **inet.3** routing table has two additional routes (to 20.20.20.2 and 30.30.30.3) with next hops derived from the same segment list path-1 with same attributes as the **to** address route.

Options

active—(Optional) Install the route into the inet.0 or inet6.0 routing table. This allows you to issue a **ping** or **traceroute** command on this address.

NOTE: The **install *destination-prefix* active** statement is not supported on static LSPs. When the **install *destination-prefix* active** statement is configured for a static LSP, the MPLS routes do not get installed into the inet.0 routing table.

destination-prefix—IPv4 or IPv6 address to associate with the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table](#) | 483

ingress-policy

Syntax

```
ingress-policy [ ingress-policy-names ];
```

Hierarchy Level

```
[edit logical-system logical-system-name protocols ldp entropy-label],  
[edit logical-system logical-system-name protocols ldp oam],  
[edit protocols ldp entropy-label],  
[edit protocols ldp oam]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced at the **[edit protocols ldp entropy-label]** hierarchy level in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Configure an LDP ingress policy for either the entropy label or Operation, Administration, and Management (OAM).

For OAM, configure the ingress policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under **[edit protocols ldp oam bfd-liveness-detection]** are applied.

Options

ingress-policy-names—Specify the names of the ingress policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring OAM Ingress Policies for LDP | 1238](#)

[Configuring the Entropy Label for LSPs | 573](#)

interface (Protocols MPLS)

Syntax

```
interface (interface-name | all) {
  disable;
  admin-group [ group-names ];
  srlg srlg-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Enable MPLS on one or more interfaces.

Options

interface-name—Name of the interface on which to configure MPLS. To configure all interfaces, specify **all**. For details about specifying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

srlg srlg-name—Name of the SRLG to associate with an interface.

The remaining options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs | 607](#)

[Example: Configuring SRLG | 278](#)

interface (MPLS)

Syntax

```
interface (all | interface-name);
```

Hierarchy Level

```
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Enable MPLS on all interfaces on the switch or on the specified interface.

Default

MPLS is disabled.

Options

all—All interfaces on the switch.

interface-name—Name of an interface:

- Aggregated Ethernet—**aex**
- Gigabit Ethernet—**ge-fpc/pic/port**

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)

[Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

inter-domain

Syntax

```
inter-domain;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path label-switched-path-name],  
[edit protocols mpls label-switched-path label-switched-path-name]
```

Release Information

Statement introduced in Junos OS Release 10.2.

Description

Allows the router to search for routes in the IGP database. You need to configure this statement on routers that might be unable to locate a path using intra-domain CSPF (by looking in the traffic engineering database (TED)). When you configure inter-area or inter-AS LSPs, the **inter-domain** statement is required.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring an LSP Across ASs | 561](#)

[label-switched-path | 2653](#)

ip-tunnel-rpf-check

Syntax

```
ip-tunnel-rpf-check {  
    mode (strict | loose);  
    fail-filter filter-name;  
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name routing-options forwarding-table]
```

Release Information

Statement introduced in Junos OS Release 17.1 for MX Series Routers with MICs.

Description

Configure the system to enable anti-spoofing protection for next-hop-based dynamic tunnels, where reverse path forwarding checks are placed to ensure that the tunnel traffic is received from a legitimate source through designated IP tunnel, where the source is reachable on the same tunnel on which the packet was received.

When a packet comes from a nondesignated source, the reverse path forwarding check fails in the strict mode, and passes in the loose mode. When a packet comes from a nonexistent source, the reverse path forwarding check fails.

By default, the reverse path forwarding check is in strict mode, where the packets are not forwarded if the source of the packet is from a nondesignated tunnel.

Options

mode (strict | loose)—(Optional) Specify the mode of the reverse path forwarding check to enable anti-spoofing protection for next-hop-based dynamic tunnels.

In the strict mode (default), the reverse path forwarding check fails when the packet is received from a nondesignated tunnel source. The check passes only for packets from designated tunnels.

In the loose mode, the reverse path forwarding check passes even if the packet is received from a nondesignated tunnel source.

When the packet is from a nonexistent tunnel source, the reverse path forwarding check fails in both the strict and loose modes.

Default: If you omit the mode statement, the default behavior is strict mode.

fail-filter *filter-name*—(Optional) Attach a filter to the Layer 3 VPN to log packets that failed the reverse path forwarding check.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview | 385](#)

[Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels | 388](#)

ipv6-tunneling

Syntax

```
ipv6-tunneling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.

Description

Allow IPv6 routes to be resolved over an MPLS network by converting LDP and RSVP routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks | 353](#)

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;
  cross-credibility-cspf;
  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 {
```



```

lsp-external-controller;
encoding-type (ethernet | packet | pdh | sonet-sdh);
gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 | pos-scrambling-crc-32 |
pos-no-scrambling-crc-32 | ppp);
signal-bandwidth type;
switching-type (fiber | lambda | psc-1 | tdm);
}
metric metric;
no-cspf;
no-decrement-ttl;
node-link-protection;
optimize-timer seconds;
p2mp lsp-name;
policing {
    filter filter-name;
    no-auto-policing;
}
preference preference;
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
}
priority setup-priority reservation-priority;
(random | least-fill | most-fill);

```



```

(record | no-record);
retry-limit number;
retry-timer seconds;
revert-timer seconds;
secondary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
}
soft-preemption;
standby;
to address;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]

```


Release Information

Statement introduced before Junos OS Release 7.4.

cross-credibility-cspf option introduced in Junos OS Release 14.2.

self-ping-duration and **no-self-ping** options introduced in Junos OS Release 16.1.

Description

Configure an LSP to use in dynamic MPLS. When configuring an LSP, you must specify the address of the egress router in the **to** statement. All remaining statements are optional.

Options

lsp-name—Name that identifies the LSP. The name can be up to 64 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique within the ingress router.

cross-credibility-cspf—Enable path computation across credibility levels. The constraint path computation is run across multi-protocol links and nodes, instead of a credibility-by-credibility basis.

link-protection—Enable protection for LSP from link faults only.

no-self-ping—Disable self-ping for the LSP.

NOTE: Starting in Junos OS 17.4, you can override the behavior of self-ping running by default using ***no-self-ping*** statement.

self-ping-duration seconds—Specify the duration (in seconds) for which to run the self-ping mechanism unless the ping succeeds sooner.

LSP self-ping for an LSP starts at the ingress label edge router (LER), once a Resv message for that LSP has been received. The configured self-ping time indicates the duration for which the self-ping mechanism runs once the LSP instance is signaled to be **UP**.

By default, self-ping is enabled. The LSP types like CCC, P2MP, VLAN-based , and non-default instances do not support self-ping .

The self-ping mechanism runs until the self-ping probe is received back (at which point the traffic is immediately switched to it) , or until the configured self-ping duration for the LSP is over (at which point traffic is switched over).

When LSP self-ping-duration is enabled, the LSP behavior reverts back to a timer-based mechanism similar to the ***optimize-switchover-delay***, where a specific amount of time is provided for all the downstream nodes to install the forwarding path before switchover. When LSP self-ping-duration is not enabled, the default behavior is to wait for an infinite amount of time for the self-ping to succeed before switching the traffic.

Range: 1 through 65535 seconds

NOTE: Starting in Junos OS 17.4R1, the default time-out duration for which the self-ping runs on an LSP instance is reduced from 65535 (runs until success) to 1800 seconds. You can also configure the self ping duration value between 1 to 65535 (runs until success) seconds.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Ingress and Egress Router Addresses for LSPs | 514](#)

[Configuring Primary and Secondary LSPs | 565](#)

label-switched-path

Syntax

```
label-switched-path lsp-name to remote-provider-edge-switch;
```

Hierarchy Level

```
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Define a label-switched path (LSP) to the remote provider edge switch to use for MPLS traffic. You must specify this statement on the provider edge switch.

Options

lsp-name —Name that identifies the LSP. The name can be up to 32 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique on the ingress switch.

remote-provider-edge-switch —Either the loopback address or the switch address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches](#) | 59

[Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\)](#) | 1318

[Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\)](#) | 1320

label-switched-path-template (Container LSP)

Syntax

```
label-switched-path-template {  
    (default-template | lsp-template-name);  
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name]  
[edit routing-instances instance-name provider-tunnel]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

Statement introduced in Junos OS Release 18.2. under the heirarchy level [edit routing-instances *instance-name* provider-tunnel]

Description

Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs.

Options

default-template—Specify that the default LSP template be used for the dynamically generated LSPs.

lsp-template-name—Specify the name of an LSP to be used as a template for the dynamically generated LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[container-label-switched-path](#) | 2579

ldp-tunneling

Syntax

```
ldp-tunneling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable the LSP to be used for LDP tunneling.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Enabling LDP over RSVP-Established LSPs](#) | 1117

least-fill

See

[random](#)

link-protection (Dynamic LSPs)

Syntax

```
link-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Enable link protection on the specified LSP, which helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails. For point-to-multipoint LSPs, including this statement extends link protection to all of the paths used by the LSP.

To fully enable link protection, you must also include the **link-protection** statement at the **[edit protocols rsvp interface *interface-name*]** or **[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]** hierarchy level.

Default

Link protection is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Link Protection for Point-to-Multipoint LSPs | 724](#)

[Configuring Node Protection or Link Protection for LSPs | 879](#)

[link-protection \(RSVP\) | 2861](#)

link-protection (Static LSPs)

Syntax

```
link-protection bypass-name name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],  
[edit protocols mpls static-label-switched-path lsp-name ingress],  
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Enable link protection on the specified static LSP. Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.

Default

Link protection is disabled.

Options

bypass-name *name*—Bypass LSP name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Static LSPs | 603](#)

[Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP | 695](#)

load-balance-label-capability

Syntax

```
load-balance-label-capability;
```

Hierarchy Level

```
[edit forwarding-options]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Description

Enables the router to push and pop the load balancing label and causes LDP and RSVP to advertise the entropy label TLV to neighboring routers.

The **load-balance-label-capability** and **no-load-balance-label-capability** statements at the **[edit forwarding-options]** hierarchy level are mutually exclusive, and at a given point in time, configuring one statement overrides the other.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Entropy Label for LSPs | 573](#)

[no-load-balance-label-capability | 2703](#)

log-updown (Protocols MPLS)

Syntax

```
log-updown {  
  no-trap {  
    mpls-lsp-traps;  
    rfc3812-traps;  
  }  
  (syslog | no-syslog);  
  trap;  
  trap-path-down;  
  trap-path-up;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

The **mpls-lsp-traps** and **rfc-3812-traps** options added in Junos OS Release 9.0.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Log a message or send an SNMP trap whenever an LSP makes a transition from up to down, or vice versa, and whenever an LSP switches from one active path to another. Only the ingress router performs these operations.

NOTE: System log messages for LSPs are generated by default. To disable the default logging of messages for LSPs, configure the **no-syslog** option under the **log-updown** statement.

Default

There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.

Options

no-syslog—Do not log a message to the system log file.

no-trap—Do not send an SNMP trap.

syslog—Log a message to the system log file.

trap—Send an SNMP trap.

trap-path-down—Send an SNMP trap when an LSP path goes down.

trap-path-up—Send an SNMP trap when an LSP path comes up.

The **no-trap** statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[System Log Messages and SNMP Traps for MPLS | 128](#)

Network Management and Monitoring Guide

[no-trap | 2707](#)

[traceoptions \(Protocols MPLS\) | 2809](#)

longest-match

Syntax

```
longest-match {  
  policy value |(expression) | [values ];  
}
```

Hierarchy Level

```
[edit protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

Description

Enable longest match to allow LDP to learn the routes aggregated or summarized across OSPF areas or IS-IS levels in inter-domain.

Options

policy *value* |(expression) | [*values*]— Specify policy to provide per prefix granularity.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Longest Match Support for LDP Overview | 933](#)

[Example: Configuring Longest Match for LDP | 940](#)

[Configuring Longest Match for LDP | 939](#)

loss (querier)

Syntax

```
loss {
  traffic-class tc-value {
    average-sample-size sample size;
    loss-threshold loss threshold value;
    loss-threshold-window number of samples for loss threshold;
    measurement-quantity bytes|packets;
    query-interval milliseconds;
  }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure loss measurement options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements | 273](#)

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

loss (responder)

Syntax

```
loss {
  min-query-interval milliseconds;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring responder],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring responder]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure loss measurement options.

Options

min-query-interval *milliseconds*—(Optional) Specify the minimum query interval that the responder supports. If the minimum query interval of the responder is greater than the query interval configured at the querier, the effective message query rate is the minimum query interval configured for the responder.

Default: 10 seconds

Range: 1000 through 4294967295 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements | 273](#)

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

loss-delay (querier)

Syntax

```
loss-delay {
  traffic-class tc-value {
    average-sample-size sample size;
    loss-threshold loss threshold value;
    loss-threshold-window number of samples for loss threshold;
    measurement-quantity bytes|packets;
    padding-size size;
    query-interval milliseconds;
    rtt-delay-threshold rtt threshold value;
    twcd-delay-threshold twcd threshold value;
  }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure combined loss-delay measurement options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements | 273](#)

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

Isp-attributes

Syntax

```

lsp-attributes {
  encoding-type (ethernet | packet | pdh | sonet-sdh);
  gpid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 | pos-scrambling-crc-32 |
    pos-no-scrambling-crc-32 | ppp);
  signal-bandwidth type;
  switching-type (fiber | lambda | psc-1 | tdm);
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]

```

Release Information

Statement introduced before Junos OS Release 7.4.

pos-scrambling-crc-16, **pos-no-scrambling-crc-16**, **pos-scrambling-crc-32**, and **pos-no-scrambling-crc-32** options added in Junos OS Release 8.0.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring MPLS LSPs for GMPLS](#) | 1372

lsping-channel-type

Syntax

```
lsping-channle-type {
  ipv4;
  on-demand-cv;
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name oam mpls-tp-mode]
[edit protocols mpls oam mpls-tp-mode]
```

Release Information

Statement introduced in Junos OS Release 16.1.

Description

Specify the control-channel types for MPLS-TP mode. By default, LSPING (0x0008) is used, and the GACH-TLV is used along with this channel type.

As per RFC 7026, GACH-TLV is deprecated for ipv4 and on-demand-cv channel types.

Options

ipv4—Channel type 0x0021. This channel type uses the IP/UDP encapsulation and provides interoperability support with other vendor devices using IP addressing.

on-demand-cv—Channel type 0x0025. This is a new pseudowire channel type and is used for on-demand CV without IP/UDP encapsulation, where IP addressing is not available or non-IP encapsulation is preferred.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[mpls-tp-mode](#) | [2690](#)

I2vpn

Syntax

```

I2vpn {
  (control-word | no-control-word);
  encapsulation-type type;
  oam {
    bfd-liveness-detection {
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
      }
      version (1 | automatic);
    }
    ping-interval seconds;
  }
  site site-name {
    community COMM;
    control-word ;
    encapsulation-type ethernet;
    ignore-encapsulation-mismatch;
    ignore-mtu-mismatch;
    interface interface-name {
      description text;
      community COMM;
      control-word ;
      encapsulation-type ethernet;
      ignore-encapsulation-mismatch;
      ignore-mtu-mismatch;
      mtu 1500;
      no-control-word;
      oam {
        bfd-liveness-detection {
          detection-time {
            threshold milliseconds;
          }
          minimum-interval milliseconds;
        }
      }
    }
  }
}

```



```

        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        version (1 | automatic);
    }
    ping-interval seconds; seconds;
}
remote-site-id remote-site-id;
target-attachment-identifier identifier;
}
mtu 1500;
no-control-word;
oam {
    bfd-liveness-detection {
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        version (1 | automatic);
    }
    ping-interval seconds; seconds;
}
site-identifier identifier;
site-preference preference-value {
    backup;
    primary;
}
source-attachment-identifier identifier;
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}

```



```
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols],  
[edit routing-instances routing-instance-name protocols]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Enable a Layer 2 VPN routing instance on a PE router or switch.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring a Layer 2 VPN Routing Instance

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

maximum-bandwidth (Protocols MPLS)

Syntax

```
maximum-bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify the maximum amount of bandwidth in bits per second (bps).

Options

bps—Maximum amount of bandwidth.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth](#) | 550

maximum-helper-recovery-time

Syntax

```
maximum-helper-recovery-time seconds;
```

Hierarchy Level

```
[edit protocols rsvp graceful-restart],  
[edit logical-systems logical-system-name protocols rsvp graceful-restart]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.

Options

seconds—Length of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.

Range: 1 through 3600

Default: 180

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for RSVP, CCC, and TCC

[maximum-helper-restart-time \(RSVP\)](#) | [2678](#)

maximum-helper-restart-time (RSVP)

Syntax

```
maximum-helper-restart-time seconds;
```

Hierarchy Level

```
[edit protocols rsvp graceful-restart],  
[edit logical-systems logical-system-name protocols rsvp graceful-restart]
```

Release Information

Statement introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

Options

seconds—The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down.

Range: 1 through 1800

Default: 60

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Graceful Restart Options for RSVP, CCC, and TCC

[maximum-helper-recovery-time](#) | [2677](#)

maximum-labels

Syntax

```
maximum-labels maximum-labels;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family mpls],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family mpls]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 switches.

Description

On the logical interface, specify the maximum number of MPLS labels upon which MPLS can operate.

Options

maximum-labels—Maximum number of labels for the protocol family.

NOTE: On PTX Series routers with third-generation FPCs, the maximum labels that can be pushed cannot exceed 8 labels.

Range: 3 through 16

Default: 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Maximum Number of MPLS Labels | 467](#)

Junos OS VPNs Library for Routing Devices

minimum-bandwidth-adjust-interval

Syntax

```
minimum-bandwidth-adjust-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify the duration (in seconds) for which minimum bandwidth is frozen.

Options

seconds—Minimum bandwidth reallocation interval, in seconds.

Range: 300 through 31,536,000 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth](#) | 550

minimum-bandwidth-adjust-threshold-change

Syntax

```
minimum-bandwidth-adjust-threshold-change percentage;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify the percentage change in maximum average bandwidth to freeze the minimum bandwidth.

Options

percentage—Percentage change in maximum average bandwidth.

Range: Range: 0 through 100 percent.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth](#) | 550

minimum-bandwidth-adjust-threshold-value

Syntax

```
minimum-bandwidth-adjust-threshold-value bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Specify the value in bits per second (bps) to freeze the minimum bandwidth if the maximum average bandwidth falls below this value.

Options

bps—Threshold value for minimum bandwidth if the maximum average bandwidth falls below the specified value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth](#) | 550

metric (Protocols MPLS)

Syntax

```
metric metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],  
[edit protocols mpls label-switched-path lsp-name],  
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Compare against another LSP or against an IGP route. To disable dynamic metric tracking, assign a fixed metric value to an LSP. If no metric is assigned, the LSP metric is dynamic and automatically tracks underlying IGP metrics.

Options

metric—LSP metric value.

Default: No metric assigned (dynamic)

Range: 1 through 16,777,215

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Static LSP Metrics](#) | 527

minimum-bandwidth

Syntax

```
minimum-bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Set the minimum bandwidth in bps for an LSP with automatic bandwidth allocation enabled.

NOTE: For a label-switched path (LSP) that has both **bandwidth** and **minimum-bandwidth** for autobandwidth configured under the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level, the LSP bandwidth is adjusted differently.

The LSP is initiated with the bandwidth value configured under the **bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level. At the expiry of the **adjust-interval** timer, the LSP bandwidth gets adjusted based on the traffic flow.

If the bandwidth to be signaled is less than the value configured under the **minimum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the minimum bandwidth.

If the bandwidth to be signaled is greater than the value configured under the **maximum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the maximum bandwidth.

Options

bps—Minimum bandwidth for the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth | 550](#)

monitor-bandwidth

Syntax

```
monitor-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Do not automatically adjust bandwidth allocation. However, the maximum average bandwidth utilization is monitored on the LSP, and the information is recorded in the MPLS statistics file.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Passive Bandwidth Utilization Monitoring | 554](#)

most-fill

See

[random](#)

mpls (Protocols)

Syntax

```
mpls { ... }
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],  
[edit protocols]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable MPLS on the router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS](#) | 55

mpls

Syntax

```
mpls {
  disable;
  class-of-service cos-value;
  no-cspf;
  no-decrement-ttl;

  advertisement-hold-time seconds;
  explicit-null;
  icmp-tunneling;
  interface (interface-name | all) {
    disable;
  }
  ipv6-tunneling;
  no-propagate-ttl;
  path path-name {
    (address | hostname) <loose | strict>;
  }
  label-switched-path lsp-name {
    disable;
    auto-bandwidth {
      adjust-interval seconds;
      adjust-threshold percentage;
      adjust-threshold-overflow-limit count;
      adjust-threshold-underflow-limit
      maximum-bandwidth bps;
      minimum-bandwidth bps;
      monitor-bandwidth;
    }
    description text-string;
    from address;
    install destination-prefix</prefix-length> <active>;
    ldp-tunneling;
    no-cspf;
    no-decrement-ttl;
    primary path-name {
      adaptive;
      select (manual | unconditional);
    }
    secondary path-name {
      adaptive;
      select (manual | unconditional);
    }
  }
}
```



```

    }
    to address;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable | no-world-readable>;
        flag flag;
    }
}
static-label-switched-path lsp-name {
    bypass bypass-name {
        description text-string;
        next-hop (address | interface-name | address/interface-name);
        to address;
    }
    ingress {
        description string;
        install {
            destination-prefix <active>;
        }
        link-protection bypass-name name;
        next-hop (address | interface-name | address/interface-name);
        to address;
    }
    transit incoming-label {
        bandwidth bps;
        description text-string;
        link-protection bypass-name name;
        next-hop (address | interface-name | address/interface-name);
        pop;
        swap out-label;
    }
    statistics {
        auto-bandwidth;
        file filename <files number> <size maximum-file-size> <world-readable | no-world-readable>;
        interval seconds;
    }
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable | no-world-readable>;
        flag flag;
    }
    traffic-engineering (bgp | bgp-igp);
}
}

```


Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Statement introduced in Release 19.4R1 for cRPD.

Description

Enable MPLS on the switch.

The remaining statements are explained separately.

Default

MPLS is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)

[Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

[Junos OS MPLS Applications Configuration Guide](#)

mpls-tp-mode

Syntax

```
mpls-tp-mode;  
lsping-channel-type;
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name oam],  
[edit protocols mpls oam]
```

Release Information

Statement introduced in Junos OS Release 12.1.

lsping-channel-type statement introduced in Junos OS Release 16.1.

Description

Enable GAL or G-Ach OAM operation without IP encapsulation on a label-switched path (LSP).

Include this statement at the **[edit protocols mpls oam]** hierarchy level to enable GAL or G-Ach OAM operation without IP encapsulation on all LSPs in the MPLS network. Include this statement at the **[edit protocols mpls label-switched-path *lsp-name* oam]** hierarchy level to enable GAL and G-Ach OAM operation without IP encapsulation on a specific LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Example: Configuring the MPLS Transport Profile for OAM](#) | 1222

mtu-signaling

Syntax

```
mtu-signaling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls path-mtu rsvp],  
[edit protocols mpls path-mtu rsvp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable MTU signaling in RSVP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Enabling MTU Signaling in RSVP](#) | 853

neighbor (Protocols Layer 2 Circuit)

Syntax

```
neighbor address {
  interface interface-name {
    backup-neighbor address {
      community name;
      hot-standby;
      psn-tunnel-endpoint address;
      standby;
      virtual-circuit-id number;
    }
    bandwidth (bandwidth | ctnumber bandwidth);
    community community-name;
    (control-word | no-control-word);
    description text;
    egress-protection {
      protected-l2circuit {
        egress-pe address;
        ingress-pe address;
        virtual-circuit-id identifier;
      }
      protector-interface interface-name;
      protector-pe address {
        context-identifier identifier;
        lsp lsp-name;
      }
    }
  }
  encapsulation-type type;
  ignore-encapsulation-mismatch;
  ignore-mtu-mismatch;
  mtu mtu-number;
  no-revert;
  protect-interface interface-name;
  pseudowire-status-tlv hot-standby-vc-on;
  psn-tunnel-endpoint address;
  revert-time seconds;
  static {
    incoming-label label;
    outgoing-label label;
    send-oam;
  }
  switchover-delay milliseconds;
}
```



```

    virtual-circuit-id identifier;
  }
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols l2circuit],
[edit protocols l2circuit]

```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Each Layer 2 circuit is represented by the logical interface connecting the local provider edge (PE) router or switch to the local customer edge (CE) router or switch. All the Layer 2 circuits using a particular remote PE router or switch designated for remote CE routers or switches are listed under the **neighbor** statement (neighbor designates the PE router or switch). Each neighbor is identified by its IP address and is usually the end-point destination for the LSP tunnel (transporting the Layer 2 circuit).

Options

address—IP address of a neighboring router or switch.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring the Neighbor Interface for the Layer 2 Circuit*

next-hop (Protocols MPLS)

Syntax

```
next-hop (address | interface-name | address/interface-name);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for IPv6 addresses in static LSP configurations added in Junos OS Release 17.2R1.

Support for IPv4 or IPv6 lookup after popping the label added in Junos OS Release 17.4R1.

Description

Location of the next hop to the destination, specified as the IPv4 or IPv6 address of the next hop, the interface name (for point-to-point interfaces only), or the ***address/interface-name*** to specify an IP address on an operational interface.

Options

address—IPv4 or IPv6 address of the next-hop router.

NOTE: IPv6 static LSPs are not supported at the [edit protocols mpls static-label-switched-path *lsp-name* ingress] hierarchy level.

interface-name—IP address of the outgoing interface. It must be a point-to-point interface. The name can be a simple or fully qualified domain name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

no-bfd-triggered-local-repair

Syntax

```
no-bfd-triggered-local-repair;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],  
[edit routing-options]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Disable Bidirectional Forwarding Detection (BFD) sessions to trigger fast reroute (FRR) using MPLS-FRR and loop-free alternates (LFAs). When this statement is configured, no BFD-triggered local repair is supported. However, logical interface down-based local repair is in force.

When using this statement to disable local repair, you also must restart routing to ensure proper behavior. To restart routing, include the **graceful-restart** command for the interior gateway protocol (IGP) used in your configuration. For example, if your IGP is OSPF, include the **graceful-restart** statement at the [edit protocols ospf] hierarchy level.

Default

BFD-triggered local repair is the default behavior. The loss of a neighbor results in BFD local repair for all next hops that derive themselves from the base next hop with which the BFD session is established.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[BFD-Triggered Local Repair for Rapid Convergence | 106](#)

[graceful-restart \(Enabling Globally\) | 2699](#)

no-cspf

Syntax

```
no-cspf;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Disable constrained-path LSP computation.

An explicit-path LSP is completely configured through operator action. Once configured, it is initiated only along the explicitly specified path.

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:

- Interior gateway protocol (IGP) topology database
- Link utilization information from extensions in the IGP link-state database
- Administrative group information from extensions in the IGP link-state database
- LSP requirements, including bandwidth, hop count, and administrative group

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.

Default

Constrained-path LSP computation enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling Constrained-Path LSP Computation](#) | 489

[Configuring Explicit-Path LSPs](#) | 686

no-decrement-ttl

Syntax

```
no-decrement-ttl;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Disable normal time-to-live (TTL) decrementing, which decrements the TTL field in the IP header by 1. This statement decrements the IP TTL by 1 before encapsulating the IP packet within an MPLS packet. When the penultimate router pops off the top label, it does not use the standard write-back procedure of writing the MPLS TTL into the IP TTL field. Therefore, the IP packet is decremented by 1. The ultimate router then decrements the packet by one more for a total cloud appearance of 2, thus hiding the network topology.

Default

Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling Normal TTL Decrementing](#) | 562

[no-propagate-ttl](#) | 2705

graceful-restart (Enabling Globally)

Syntax

```
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;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],  
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options],  
[edit routing-options],  
[edit routing-instances routing-instance-name routing-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.1 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

You configure the graceful restart routing option globally to enable the feature, but not to enable graceful restart for all routing protocols in a routing instance. To enable graceful restart globally, include the graceful-restart statement under the **[edit routing options]** hierarchy level. This enables graceful restart globally for all routing protocols. 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.

Default

Graceful restart is disabled by default.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Enabling Graceful Restart

Configuring Routing Protocols Graceful Restart

Configuring Graceful Restart for MPLS-Related Protocols

Configuring VPN Graceful Restart

Configuring Logical System Graceful Restart

Configuring Graceful Restart for QFabric Systems

helper-disable (Multiple Protocols)

Syntax

```
helper-disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ldp | ospf | ospf3)  
  graceful-restart],  
[edit protocols (isis | ldp | ospf | ospf3 | rsvp) graceful-restart],  
[edit routing-instances routing-instance-name protocols (ldp | ospf | ospf3) graceful-restart]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.

Default

Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and RSVP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Routing Protocols Graceful Restart](#)

[Configuring Graceful Restart for MPLS-Related Protocols](#)

no-install-to-address

Syntax

```
no-install-to-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],  
[edit protocols mpls label-switched-path lsp-name],  
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Release Information

Statement introduced in Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Prevent the egress router address configured using the **to** statement from being installed into the inet.3 and inet.0 routing tables.

Default

The egress router address for an LSP is installed into the inet.3 and inet.0 routing tables.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Preventing the Addition of Egress Router Addresses to Routing Tables | 515](#)
[to | 2808](#)

no-load-balance-label-capability

Syntax

```
no-load-balance-label-capability;
```

Hierarchy Level

```
[edit forwarding-options]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Description

Disables advertisement of entropy label capability in LDP and RSVP.

The **load-balance-label-capability** and **no-load-balance-label-capability** statements at the **[edit forwarding-options]** hierarchy level are mutually exclusive, and at a given point in time, configuring one statement overrides the other.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[load-balance-label-capability](#) | 2664

[entropy-label](#) | 2602

[Configuring the Entropy Label for LSPs](#) | 573

no-mcast-replication

Syntax

```
no-mcast-replication;
```

Hierarchy Level

```
[edit chassis fpc slot-number pic pic-number],  
[edit chassis lcc number fpc slot-number pic pic-number]
```

Release Information

Statement introduced in Junos OS Release 11.3.

Description

For point-to-multipoint LSPs configured on T Series routers, protect the Packet Forwarding Engine (PFE) from bandwidth saturation. When a PFE does not need to replicate traffic, the PFE's bandwidth is less likely to become saturated. When you include the **no-mcast-replication** statement, the PFE is forced to be a leaf node in the binary tree. Leaf nodes, unlike branch nodes, do not replicate traffic in the process of forwarding traffic. Because leaf nodes have no children, they do not need to replicate traffic, and thus are less likely to become saturated with traffic.

Default

If you omit the **no-mcast-replication** statement, the PFE can become a branch node or a leaf node. When the PFE becomes a branch node, the PFE must replicate traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Point-to-Multipoint LSPs Overview](#) | 691

no-propagate-ttl

Syntax

```
no-propagate-ttl;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

When you add the **no-propagate-ttl** statement to the configuration or delete it from the configuration, the effect takes place immediately. There is no need to clear existing RSVP LSPs or LDP sessions.

Default

Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling Normal TTL Decrementing | 562](#)

Example: Diagnosing Networking Problems Related to Layer 3 VPNs by Disabling TTL Decrementing (on Layer 3 VPNs User Guide for Routing Devices or in the Junos VPNs Configuration Guide)

[no-decrement-ttl | 2698](#)

no-transit-statistics

Syntax

```
no-transit-statistics;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls statistics],  
[edit protocols mpls statistics]
```

Release Information

Statement introduced in Junos OS Release 10.2 for PTX Series.

Description

(PTX Series only) Disables the collection of MPLS statistics for LSPs transiting the router.

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 | 242](#)

[statistics | 2799](#)

no-trap

Syntax

```
no-trap {  
    mpls-lsp-traps;  
    rfc-3812-traps;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls log-updown],  
[edit protocols mpls log-updown]
```

Release Information

Statement introduced before Junos OS Release 7.4.

The **mpls-lsp-traps** and **rfc-3812-traps** options added in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Prevent the transmission of SNMP traps.

Options

mpls-lsp-traps—Block the MPLS LSP traps defined in the **rfc-3812-traps**, but allows the **rfc3812.mib** traps.

rfc-3812-traps—Block the traps defined in the **rfc3812.mib**, but allows the MPLS LSP traps defined in the **jnx-mpls.mib**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[System Log Messages and SNMP Traps for MPLS | 128](#)

Network Management and Monitoring Guide

[traceoptions \(Protocols MPLS\) | 2809](#)

node-protection (Static LSP)

Syntax

```
node-protection bypass-name name next-next-label label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],  
[edit protocols mpls static-label-switched-path lsp-name ingress],  
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Release Information

Statement introduced in JUNOS Release 10.1.

Description

Enable node protection on the specified static bypass LSP. Node protection ensures that traffic from an LSP traversing a neighboring router can continue to reach its destination even if the neighboring router fails.

Default

Node protection is disabled.

Options

bypass-name *name*—Bypass LSP name.

next-next-label *label*—Bypass LSP name.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Static LSPs](#) | 603

normalization

Syntax

```
normalization {
    failover-normalization;
    no-incremental-normalize;
    normalization-retry-duration seconds;
    normalization-retry-limits number;
    normalize-interval seconds;
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name splitting-merging]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced in Junos OS Release 17.2R1 for QFX Series switches.

Description

Perform normalization.

Options

failover-normalization—Enable the ingress router to pro-actively normalize or re-distribute traffic when a link or node failure happens on a member LSP. A member LSP can go down between two scheduled normalization events because of a link-failure or pre-emption.

Default: Disabled

no-incremental-normalize—Disables automatic switchover by the ingress router to a new instance of the container LSP until the desired demand is satisfied, although the given number of LSPs can be successfully signaled such that the new aggregate bandwidth value exceeds the old aggregate bandwidth value.

Default: False (disabled)

normalization-retry-duration *seconds*—Specifies the duration before which the ingress router performs a normalization reattempt when the previous normalization has not been successful. Normalization is done until a sufficient number of LSPs come up with an aggregate bandwidth that is more than the current aggregate or desired bandwidth.

Default: 30 seconds

normalization-retry-limits *number*—Specifies the maximum number of times the ingress router performs normalization reattempts until a sufficient number of LSPs come up successfully with new bandwidth values.

Default: 1

normalize-interval *seconds*—Specifies the duration between two normalization events.

Range: 21600 seconds through 6 hours

Default: 21600 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [splitting-merging](#) | [2788](#)

oam (Protocols MPLS)

Syntax

```

oam {
  bfd-liveness-detection{
    failure-action teardown;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
  }
  lsp-ping-interval seconds;
  mpls-tp-mode;
  performance-monitoring {
    querier {
      loss {
        traffic-class tc-value {
          query-interval milliseconds;
          measurement-quantity bytes|packets;
          average-sample-size sample size;
          loss-threshold loss threshold value;
          loss-threshold-window number of samples for loss threshold;
        }
      }
    }
    delay {
      traffic-class tc-value {
        query-interval milliseconds;
        padding-size size;
        average-sample-size sample size;
        rtt-delay-threshold rtt threshold value;
        twcd-delay-threshold twcd threshold value;
      }
    }
  }
  loss-delay {
    traffic-class tc-value {
      query-interval milliseconds;
      measurement-quantity bytes|packets;
      padding-size size;
      average-sample-size sample size;
      loss-threshold loss threshold value;
      loss-threshold-window number of samples for loss threshold;
      rtt-delay-threshold rtt threshold value;
      twcd-delay-threshold twcd threshold value;
    }
  }
}

```



```

    }
  }
  responder {
    loss {
      min-query-interval milliseconds;
    }
    delay {
      min-query-interval milliseconds;
    }
  }
}
}
}

```

Hierarchy Level

```

[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
[edit protocols mpls label-switched-path lsp-name primary path-name]

```

Release Information

Statement introduced in Junos OS Release 7.6.

lsp-ping-interval option introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

performance-monitoring configuration statement introduced in Junos OS Release 15.1.

Description

Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.

Options

lsp-ping-interval *seconds*—Specify the duration of the LSP ping interval in seconds. To issue a ping on an RSVP-signaled LSP, use the **ping mpls rsvp** command.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring BFD for MPLS IPv4 LSPs | 109](#)

optimize-adaptive-teardown

Syntax

```
optimize-adaptive-teardown {  
  p2p:  
    delay value (3..65535 seconds)  
}
```

Hierarchy Level

```
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 15.1R1.

Description

Make use of a new feedback mechanism from TAG library which relies on RPD infrastructure to decide when all the routes using the old LSP instance have fully shifted to the new LSP instance after MBB switchover. When this statement is configured, the **optimize-hold-dead-delay** statement, which delays the teardown of the old LSP instance after MBB switchover, is ignored.

Options

p2p—Only point-to-point LSPs configured in the system will be affected.

delay—Delays the tearing down of old optimized LSP paths based on the configured value. The value configured is in the range of 3..65535 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Optimizing Signaled LSPs | 541](#)[Achieving a Make-Before-Break, Hitless Switchover for LSPs | 538](#)

optimize-aggressive

Syntax

```
optimize-aggressive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

If enabled, the LSP reoptimization is based solely on the IGP metric. The reoptimization process ignores the available bandwidth ratio calculations, the least-fill 10 percent congestion improvement rule, and the hop-counts rule. This statement makes reoptimization more aggressive than the default.

Default

Aggressive optimization is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Optimizing Signaled LSPs](#) | 541

optimize-hold-dead-delay

Syntax

```
optimized-hold-dead-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switch-path lsp-name],
[edit protocols mpls],
[edit protocols mpls label-switch-path lsp-name]
```

Description

Allows you to specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths. This delay timer starts when the timer specified by the **optimize-switchover-delay** statement has elapsed, which is typically 30 seconds, and at the start of the next retry sequence (in other words, the delay is not an absolute countdown of the seconds configured here).

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.

Options

seconds—Configure the time in seconds to wait before tearing down the old paths that were in use prior to the last LSP optimization.

Default: 60 to 90 seconds

Range: 0 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Optimizing Signaled LSPs | 541](#)

[optimize-switchover-delay | 2716](#)

[optimize-timer | 2717](#)

optimize-switchover-delay

Syntax

```
optimize-switchover-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced in Junos OS Release 11.1R1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Delays the switch over of LSPs to newly optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths.

Options

seconds—Configure the time in seconds to wait before switching LSPs to newly optimized paths.

Default: 1 second

Range: 1 through 900 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Optimizing Signaled LSPs | 541](#)

[optimize-hold-dead-delay | 2715](#)

[optimize-timer | 2717](#)

optimize-timer (Protocols MPLS)

Syntax

```
optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Enable periodic reoptimization of an LSP that is already set up. If topology changes occur, an existing path might become suboptimal, and a subsequent recomputation might be able to determine a better path. This feature is useful only on LSPs for which constrained-path computation is enabled; that is, for which the **no-cspf** statement is not configured. Also, you only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers).

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.

Default

The optimize timer is disabled.

Options

seconds—Length of the optimize timer, in seconds.

Range: 0 through 65,535 seconds

Default: 0 seconds (the optimize timer is disabled)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Optimizing Signaled LSPs](#) | 541

p2mp (Protocols MPLS)

Syntax

```
p2mp p2mp-lsp-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description

Specify an LSP as either a point-to-multipoint LSP or as a branch LSP of a point-to-multipoint LSP by specifying the point-to-multipoint LSP path name.

Options

p2mp-lsp-name—Name of the point-to-multipoint LSP path that identifies the sequence of nodes that form the point-to-multipoint LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Primary Point-to-Multipoint LSP](#) | 721

p2mp-lsp-next-hop

Syntax

```
p2mp-lsp-next-hop {
    metric metric;
    preference preference;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options static route
destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix],
[edit routing-options static route destination-prefix]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify a point-to-multipoint LSP as the next hop for a static route, and configure an independent metric or preference on that next-hop LSP.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Static Unicast Routes for Point-to-Multipoint LSPs | 611](#)

[Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP | 695](#)

[Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems](#)

path (Protocols MPLS)

Syntax

```
path path-name {
    abstract-hop-name (abstract | loose | loose-link | strict);
    (address | hostname) <strict | loose>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

abstract-hop-name option introduced in Junos OS Release 17.1 for all platforms.

Description

Create a named path and optionally specify the sequence of explicit routers that form the path.

You must include this statement when configuring explicit LSPs.

Options

abstract-hop-name—Name of the predefined abstract hop. The abstract hop can be used in combination with real IP next hops. An abstract hop is traversed by traversing the member nodes. This traversal can be done by either links that satisfy the logical combination of defined constituent attributes, or by any kind of link. This choice is controlled by the use of abstract hop qualifiers – abstract, loose, loose-link, and strict.

abstract—Indicate that the next hop configured in the **path** statement is an abstract hop..

loose-link—Indicate that the next hop in the **path** statement is a loose-link abstract hop. This means that the LSP cannot traverse other routers before reaching this router. In other words, the abstract hop of type loose-link is processed only if any of the viable routers is reached in constraint through a link of associated abstract hop membership.

loose—Indicate that the next hop in the **path** statement is a loose abstract hop. The path can traverse any real nodes that do not have abstract hop membership, before reaching a node with abstract hop membership, which is a feasible starting point for processing the next abstract hop.

strict—Indicate that the next hop in the **path** statement is a strict abstract hop. After the last processed hop in the constraint list, the path can traverse any real nodes that do not have abstract hop membership, before reaching a node with abstract hop membership, which is a feasible starting point for processing the next abstract hop.

address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is **loose**. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.

Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.

hostname—See **address**.

Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.

loose—(Optional) Indicate that the next address in the **path** statement is a loose link. This means that the LSP can traverse through other routers before reaching this router.

Default: **strict**

path-name—Name that identifies the sequence of nodes that form an LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.

strict—(Optional) Indicate that the LSP must go to the next address specified in the **path** statement without traversing other nodes. This is the default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Creating Named Paths | 516](#)

[abstract-hop | 2534](#)

path

Syntax

```
path destination {
  <address | hostname> <strict | loose>
}
```

Hierarchy Level

[edit protocols [mpls](#)]

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Configure path protection on your MPLS network.

Options

destination —Name of a label switched path (LSP). In addition to specifying the name of the configured LSP, you can include some other designation such as **primary-path**.

address —(Optional) IP address of each transit switch (or the IP address of the loopback interface on the switch) in the LSP. If you want to control exactly which switches are selected for the LSP, specify the address or hostname of each transit switch. Specify the addresses in order, starting with the first provider (transit) switch, and continuing sequentially along the path until reaching the egress provider edge switch.

Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.

hostname —(Optional) See **address** .

Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.

loose—(Optional) Indicates that the next address in the **path** statement is a loose link. This means that the LSP can traverse through other switches before reaching this switch.

Default: **strict**

strict—(Optional) Indicates that the LSP must go to the next address specified in the **path** statement without traversing other switches. This is the default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Path Protection in an MPLS Network \(CLI Procedure\) | 145](#)

path-mtu

Syntax

```
path-mtu {  
    allow-fragmentation;  
    rsvp {  
        mtu-signaling;  
    }  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Configure MTU options for MPLS paths, including packet fragmentation and MTU signaling.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring MTU Signaling in RSVP | 853](#)

per-prefix-label

Syntax

```
per-prefix-label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit logical-systems logical-system-name protocols bgp group group-name family inet labeled-unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp group group-name family inet labeled-unicast],
[edit routing-instances instance-name logical-systems logical-system-name protocols bgp family inet labeled-unicast],
[edit routing-instances instance-name logical-systems logical-system-name protocols bgp group group-name family
  inet labeled-unicast],
[edit routing-instances instance-name protocols bgp family inet labeled-unicast],
[edit routing-instances instance-name protocols bgp group group-name family inet labeled-unicast]
```

Release Information

Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.

Statement introduced in Junos OS Release 12.3 for M Series, T Series, and MX Series routers.

Description

Allocate a unique label for each prefix. The **per-prefix-label** statement helps minimize packet loss in most deployments.

Although allocating a label for each prefix is not generally ideal for scaling, it is assumed that a small number of labels are used for BGP labeled-unicast. When labeled BGP is used to set up transport label-switched paths (LSPs), the common case is that each prefix has a unique next hop. Thus, the use of per-prefix labels does not have an adverse scaling impact. On the contrary, the use of per-prefix labels reduces churn in the network when multipath load balancing is enabled for IPv4 labeled-unicast, and a subset of the paths are withdrawn for some reason.

The advantage of per-prefix labeling is that the advertised upstream label is more stable during network changes. That is, if the downstream label changes, the advertised upstream label remains the same under most scenarios. This way, the upstream router is isolated from the downstream network change, and the overall network is more stable. The greater stability of the advertised upstream label helps to reduce traffic loss during many different network change scenarios.

Default

By default, label allocation is per next-hop router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [MPLS Label Allocation](#) | 420

performance-monitoring (Protocols MPLS)

Syntax

```
performance-monitoring {
  querier {
    delay {
      traffic-class tc-value {
        average-sample-size sample size;
        padding-size size;
        query-interval milliseconds;
        rtt-delay-threshold rtt threshold value;
        twcd-delay-threshold twcd threshold value;
      }
    }
    loss {
      traffic-class tc-value {
        average-sample-size sample size;
        loss-threshold loss threshold value;
        loss-threshold-window number of samples for loss threshold;
        measurement-quantity bytes|packets;
        query-interval milliseconds;
      }
    }
    loss-delay {
      traffic-class tc-value {
        average-sample-size sample size;
        loss-threshold loss threshold value;
        loss-threshold-window number of samples for loss threshold;
        measurement-quantity bytes|packets;
        padding-size size;
        query-interval milliseconds;
        rtt-delay-threshold rtt threshold value;
        twcd-delay-threshold twcd threshold value;
      }
    }
  }
  responder {
    delay {
      min-query-interval milliseconds;
    }
    loss {
      min-query-interval milliseconds;
    }
  }
}
```



```
}
```

Hierarchy Level

```
[edit protocols mpls oam],  
[edit protocols mpls label-switched-path lsp-name oam],  
[edit protocols mpls label-switched-path lsp-name primary path-name oam],  
[edit protocols mpls label-switched-path lsp-name secondary path-name oam]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure performance monitoring options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

policing (Protocols MPLS)

Syntax

```
policing {  
    filter filter-name;  
    no-auto-policing;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],  
[edit protocols mpls label-switched-path lsp-name],  
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the policing filter for the LSP.

Options

filter *filter-name*—Specify the name of the policing filter.

no-auto-policing—Disable automatic policing on this LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Policers for LSPs | 121](#)

[auto-policing | 2562](#)

policing

Syntax

```
policing (filter filter-name | no-automatic-policing);
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name]  
[edit interfaces interface-id unit number-of-logical-unit family inet address ip-address]
```

Release Information

Statement introduced in Junos OS Release 10.1 for EX Series switches.

Description

Apply a rate-limiting policer as the specified policing filter:

- To the LSP for MPLS over CCC.
- To the customer-edge interface for IP over MPLS.

Options

filter *filter-name*—Specify the name of the policing filter.

no-automatic-policing—Disable automatic policing on this LSP.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

policer

Configuring Policers to Control Traffic Rates (CLI Procedure)

[Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)

[Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)

policy-multipath

Syntax

```
policy-multipath policy [ policy ]
  traceoptions <file filename <files files> <size size> <(world-readable | no-world-readable)>> name detail disable
  receive send
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options rib],
[edit logical-systems name routing-options rib],
[edit logical-systems name tenants name routing-instances name routing-options rib],
[edit routing-instances name routing-options rib],
[edit routing-options rib],
[edit tenants name routing-instances name routing-options rib]
```

Release Information

Statement introduced in Junos OS Release 19.1R1 for all platforms.

Description

Create policy-based multipath route using a combination of segment routing traffic-engineered (SR-TE) LDP or RSVP routes and SR-TE IP routes. You can resolve BGP service routes over the multipath route, and apply export policies to steer traffic differently for different prefixes. The policy-based multipath feature is supported for both IP and IPv6 protocols.

Options

policy—Import policy to create policy-based multipath.

NOTE: This statement is supported only at the **[edit routing-option policy-multipath]** hierarchy level.

Any action commands configured in the policy, such as apply, is evaluated using the active route. For non-active routes, the policy is applied to check if the routes can participate in the multipath route or not. Multipath routes inherit all attributes of the active route. These attributes can be modified using the multipath policy configuration.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing

RELATED DOCUMENTATION

| [Policy-Based Multipath Routes Overview](#) | 409

policy-statement

Syntax

```

policy-statement policy-name {
  term term-name {
    from {
      as-path-unique-count count (equal | orhigher | orlower);
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      protocol protocol-name;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
      tag value;
      traffic-engineering;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
  then {
    aggregate-bandwidth;
    dynamic-tunnel-attributes dynamic-tunnel-attributes;
    limit-bandwidth limit-bandwidth;
    multipath-resolve multipath-resolve;
    no-entropy-label-capability;
    prefix-segment {
      index index;
      node-segment;
    }
    priority (high | medium | low);
    resolution-map map-name;
  }
}

```

Hierarchy Level

```

[edit dynamic-profiles profile-name policy-options],
[edit logical-systems logical-system-name policy-options],

```


[edit policy-options]

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5.

Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.

inet-mdt option introduced in Junos OS Release 10.0R2.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

route-target option introduced in Junos OS Release 12.2.

Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.

protocol and **traffic-engineering** options introduced in Junos OS Release 14.2.

no-entropy-label-capability option introduced in Junos OS Release 15.1.

priority and **tag value** options introduced in Junos OS Release 17.1.

as-path-unique-count option introduced in Junos OS Release 17.2R1.

prefix-segment option introduced in Junos OS Release 17.2R1 for MX Series routers, PTX Series routers, QFX5100 switches, and QFX10000 switches.

multipath-resolve and **dynamic-tunnel-attributes** options introduced in Junos OS Release 17.3R1.

aggregate-bandwidth and **limit-bandwidth** *limit-bandwidth* options introduced in Junos OS Release 17.4R1 for MX Series, PTX Series, and QFX Series.

l-isis and *l-ospf* keywords at the **protocol** option is introduced in Junos OS Release 19.1R1.

resolution-map statement introduced in Junos OS Release 19.2R1-S1 on MX and PTX Series routers.

lsp and **lsp-regex** options introduced in Junos OS Release 19.4R1.

Description

Define a routing policy, including subroutine policies.

A *term* 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.

Each term contains a set of match conditions and a set of actions:

- 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.

The statements are explained separately.

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.

as-path-unique-count *count* (*equal* | *orhigher* | *orlower*)—(Optional) Specify a number from 0 through 1024 to filter routes based on the number of unique autonomous systems (ASs) in the AS path. Specify the match condition for the unique AS path count.

aggregate-bandwidth—(Optional) Enable BGP to advertise aggregate outbound link bandwidth for load balancing.

dynamic-tunnel-attributes *dynamic-tunnel-attributes*—(Optional) Choose a set of defined dynamic tunnel attributes for forwarding traffic over V4oV6 tunnels.

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*.

multipath-resolve *multipath-resolve*—(Optional) Enable the use of all paths for resolution over the specified prefix.

limit-bandwidth *limit-bandwidth*—(Optional) Specify the limit for advertised aggregate outbound link bandwidth for load balancing.

Range: 0 through 4,294,967,295 bytes

no-entropy-label-capability—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

priority (*high* | *medium* | *low*)—(Optional) Configure the priority for an IS-IS route to change the default order in which the routes are installed in the routing table, in the event of a network topology change.

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, 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.

Starting in Junos OS Release 19.1R1, you can specify options to match label IS-IS and label OSPF routes using the **l-isis** and **l-ospf** options, respectively. The **isis** options matches all IS-IS routes, excluding labelled IS-IS routes. The **ospf** option matches all OSPF routes, including OSPFv2, OSPFv3 and labelled OSPF routes.

resolution-map—(Optional) Set resolution map modes. A given resolution-map can be shared across multiple policy-statements.

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.

tag value—(Optional) A numeric value that identifies a route. You can tag certain routes to prioritize them over other routes. In the event of a network topology change, Junos OS updates these routes in the routing table before updating other routes with lower priority. You can also tag some routes to identify and reject them based on your requirement.

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*.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

dynamic-db

Understanding Source Packet Routing in Networking (SPRING)

pop

Syntax

```
pop;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],  
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Remove the label from the top of the label stack. If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs | 607](#)

[swap | 2801](#)

pop-and-forward (Protocols MPLS)

Syntax

```
pop-and-forward;
```

Hierarchy Level

```
[edit logical-systems name protocols mpls label-switched-path ],  
[edit logical-systems name routing-instances name protocols mpls label-switched-path ],  
[edit protocols mpls label-switched-path ],  
[edit routing-instances name protocols mpls label-switched-path ]
```

Release Information

Statement introduced in Junos OS Release 18.1R1 on MX Series routers, PTX Series routers, and vMX.

Description

Enable LSP as pop-and-forward with auto-delegation signaling enabled by default.

The LSP undergoes a make-before-break from a regular point-to-point LSP to a pop-and-forward LSP.

Required Privilege Level

routing

RELATED DOCUMENTATION

[Pop-and-Forward LSP Configuration](#) | 797

[show rsvp pop-and-forward](#) | 3387

[pop-and-forward \(Protocols RSVP\)](#) | 2881

preference (Protocols MPLS)

Syntax

```
preference preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Preference for the route.

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.

Options

preference—Preference to assign to the route. A route with a lower preference value is preferred.

Range: 1 through 255

Default: 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Preference Values for LSPs | 537

Configuring the Ingress Router for Static LSPs | 604

Configuring the Intermediate (Transit) and Egress Routers for Static LSPs | 607

primary (Protocols MPLS)

Syntax

```
primary path-name {
  adaptive;
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  bandwidth 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;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the primary path to use for an LSP. You can configure only one primary path.

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 lsp-name]` hierarchy level).

Options

path-name—Name of a path that you created with the **path** statement.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Primary and Secondary LSPs](#) | 565

primary**Syntax**

```
primary path-name;
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Specify the primary path to use for a label switched path (LSP). You can configure only one primary path.

Options

path-name —Name of the primary path that you created with the **path** statement.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Path Protection in an MPLS Network \(CLI Procedure\)](#) | 145

priority (Protocols MPLS)

Syntax

```
priority setup-priority reservation-priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the setup priority and reservation priority for an LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. Sessions with lower hold priorities are preempted.

Options

reservation-priority—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.

Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.

Default: 0 (Once the session is set up, no other session can preempt it.)

setup-priority—Setup priority.

Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.

Default: 7 (The session cannot preempt any existing sessions.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

protection-revert-time

Syntax

```
protection-revert-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls interface interface-name static],  
[edit protocols mpls interface interface-name static]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Description

Specify the amount of time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path.

If you have configured a value of 0 seconds for the **protection-revert-time** statement and traffic is switched to the bypass path, the traffic remains on that path indefinitely. It is never switched back to the original path unless the bypass path is down or you intervene.

Options

seconds—Time in seconds.

Range: 0 through 65,535 seconds

Default: 5 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

push

Syntax

```
push out-label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name bypass],  
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],  
[edit protocols mpls static-label-switched-path lsp-name bypass],  
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Add a new label to the top of the label stack. This statement is used to configure static LSPs at ingress routers and to configure bypass LSPs for static LSPs.

Options

out-label—Manually assigned outgoing label value.

Range: 0 through 1,048,575.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pop | 2739](#)

[swap | 2801](#)

[Configuring the Ingress Router for Static LSPs | 604](#)

random

Syntax

```
(random | least-fill | most-fill);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the preferred path when several equal-cost candidate paths to a destination exist, and prefer the path with the highest available bandwidth (with the largest minimum available bandwidth ratio). The available bandwidth ratio of a link is the available bandwidth on a link divided by the maximum reservable bandwidth on the link.

- **least-fill**—Prefer the path with the most available bandwidth (with the largest available bandwidth ratio).
- **most-fill**—Prefer the path with the least available bandwidth (with the minimum available bandwidth ratio). The minimum available bandwidth ratio of a path is the smallest available bandwidth ratio belonging to any of the links in the path.
- **random**—Choose the path at random.

Default

random

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring CSPF Tie Breaking](#) | 489

record

Syntax

```
(record | no-record);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path (primary | secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify whether an LSP should actively record the routes in the path. Recording routes requires that all transit routers support the RSVP Record Route object. Recording routes can be useful for diagnostics and loop detection.

Default

Record routes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling Path Route Recording by LSPs](#) | 537

remote-interface-switch

Syntax

```
remote-interface-switch connection-name {
  interface interface-name.unit-number;
  receive-lsp label-switched-path;
  transmit-lsp label-switched-path;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],
[edit protocols connections \(MPLS\)]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Configure MPLS LSP tunnel cross-connects. This makes an association between a CCC interface and two LSPs, one for transmitting MPLS packets from the local provider edge switch to the remote provider edge switch and the other for receiving MPLS packets on the local provider edge switch from the remote provider edge switch.

Options

connection-name—Connection name.

interface interface-name.unit-number—Interface name. Include the logical portion of the name, which corresponds to the logical unit number of the CCC interface.

receive-lsp label-switched-path—Name of the LSP from the connection's source. This LSP name was specified by the **label-switched-path** statement on the remote provider edge switch in the **protocols mpls** stanza.

transmit-lsp label-switched-path—Name of the LSP to the connection's destination. This LSP name was specified by the **label-switched-path** statement on the local provider edge switch in the **protocols mpls** stanza.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS LSP Tunnel Cross-Connects Using CCC | 1665](#)

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

MPLS Applications User Guide

remote-site-id

Syntax

```
remote-site-id remote-site-ID;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn site site-name
  interface interface-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Control the remote interface to which the interface should connect. If you do not explicitly configure the remote site ID, the order of the interfaces configured for the site determines the default value. This statement is optional.

Options

remote-site-ID—Identifier specifying the interface on the remote PE router the Layer 2 VPN routing instance connects to.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Remote Site ID

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

retry-limit

Syntax

```
retry-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name],
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Maximum number of times the ingress router tries to establish the primary path. This counter is reset each time a primary path is created successfully. When the limit is exceeded, no more connection attempts are made. Intervention is then required to restart the connection.

Options

number—Maximum number of tries to establish the primary path.

Range: 0 through 10,000

Default: 0 (The ingress node never stops trying to establish the primary path.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Connection Between Ingress and Egress Routers](#) | 522

retry-timer

Syntax

```
retry-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Amount of time the ingress router waits between attempts to establish the primary path.

Options

seconds—Amount of time between attempts to connect to the primary path.

Range: 1 through 600 seconds

Default: 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Connection Between Ingress and Egress Routers](#) | 522

revert-timer

Syntax

```
revert-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

BFD behavior modified in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the amount of time (in seconds) that an LSP must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted.

If you have configured BFD on the LSP, the Junos OS waits until the BFD session is restored before starting the revert timer counter.

If you have configured a value of 0 seconds for the **revert-timer** statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.

Options

seconds—Time in seconds.

Range: 0 through 65,535 seconds

Default: 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Revert Timer for LSPs](#) | 567

revert-timer

Syntax

```
revert-timer seconds;
```

Hierarchy Level

```
[edit protocols mpls],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Specify the amount of time that a label switched path (LSP) must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted.

If you have configured a value of 0 seconds for the **revert-timer** statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.

Default

60 seconds

Options

seconds —Value in seconds.

Range: 0 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Path Protection in an MPLS Network \(CLI Procedure\)](#) | 145

resignal-minimum-bandwidth

Syntax

```
resignal-minimum-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],  
[edit protocols mpls label-switched-path lsp-name auto-bandwidth]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Description

Resignal the LSP using the configured minimum bandwidth when an LSP comes back up after going down.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Automatic Bandwidth Allocation for LSPs](#) | 547

resolution-map

Syntax

```
resolution-map name {
    mode (color-only | ip-color);
}
```

Hierarchy Level

```
[edit logical-systems name policy-options],
[edit policy-options]
```

Release Information

Statement introduced in Junos OS Release 19.2R1-S1 on MX Series and PTX Series routers.

Description

Define a set of protocol next hop resolution modes.

A resolution-map can be referred by a new resolution-map action of a policy statement, which is in turn applied to a VPN service through the routing-instance **import-vrf**. A given resolution-map may be shared by multiple policy-statements.

Options

name—Resolution Map name.

mode—List of resolution modes in order that defines fallback mechanism.

Values:

- color-only—Color-only protocol next hop resolution mode.
- ip-color—Colored-IP protocol next hop resolution mode.

Required Privilege Level

routing

RELATED DOCUMENTATION

[Color-Based Mapping of VPN Services Overview | 622](#)

Static Segment Routing LSP in MPLS Network

responder (performance-monitoring)

Syntax

```
responder {
  delay {
    min-query-interval milliseconds;
  }
  loss {
    min-query-interval milliseconds;
  }
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure responder options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements](#) | 273

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview](#) | 243

[performance-monitoring \(Protocols MPLS\)](#) | 2727

rpf-check-policy (Routing Options)

Syntax

```
rpf-check-policy policy;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options multicast],  
[edit routing-options multicast]
```

Release Information

Statement introduced in Junos OS Release 8.0.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Enable you to control whether a reverse path forwarding (RPF) check is performed for a source and group entry before installing a route in the multicast forwarding cache. This makes it possible to use point-to-multipoint LSPs to distribute multicast traffic to Protocol Independent Multicast (PIM) islands situated downstream from the egress routers of the point-to-multipoint LSPs.

Options

policy—Name of the RPF check routing policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs](#) | 725

rsvp-error-hold-time

Syntax

```
rsvp-error-hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

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.

Options

seconds—Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations.

Range: 0 through 240 seconds

Default: 25 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages](#) | 1189

sampling (Protocols MPLS)

Syntax

```
sampling {
  cut-off-threshold percentile;
  use-average-aggregate;
  use-percentile percentile;
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name splitting-merging]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced in Junos OS Release 17.2R1 for the QFX Series switches.

Description

Configure traffic sampling. By default, sampling mode is set to **Max Aggregate** which means the system will compare each new sample with the previous sample. If the new sample is higher than the old sample, then the newer sample is considered **Sampled Aggregate bandwidth**

Max Aggregate Example

If normalization happens every 20min (T0, T20, T40..) then if at time T0 the traffic rate is 185Gbps, and subsequently drops to 7.5Gbps at time T3, the max-aggregate sample for the current normalization window T0-T20 will be 185Gbps. When the current normalization window expires at time T20, we use the previous sampled max-aggregate of 185Gbps to calculate the split/merge activities of the next (or now current) normalization window between T20 – T40. If traffic remains at 7.5Gbps for this second normalization period, then at time T40 the max-aggregate sample of 7.5Gbps would then be used for split/merge activities. Even though traffic volumes dropped at time T3, LSP split/merge activities would not occur until time T40 which might be unexpected. This default behavior can be modified with use-average-aggregate or use-percentile to achieve alternative normalization behavior if desired.

Options

cut-off-threshold *percentile*—Specify the percentile value to be used as a cut-off threshold in removing outlier bandwidth samples. All the aggregate bandwidth samples determined as outliers are used for computing aggregate bandwidth used at the time of normalization.

Default: 0 percentile (the ingress considers all aggregate bandwidth samples for normalization.)

Range: 0 through 100

use-average-aggregate—Specify the ingress router to take average of the aggregate samples for normalization.

This option is mutually exclusive with the **use-percentile** configuration option.

use-percentile *percentile*—Specify the ingress router to compute and use the **pth** percentile from all the bandwidth samples, and use that for normalization.

This option is mutually exclusive with the **use-average-aggregate** configuration option.

Range: 0 through 100

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [splitting-merging](#) | [2788](#)

sbfd

Syntax

```
sbfd {  
    remote-discriminator remote-discriminator;  
}
```

Hierarchy Level

[edit protocols source-packet-routing [source-routing-path](#) name primary name [bfd-liveness-detection \(LSP\)](#)],
[edit protocols source-packet-routing [source-routing-path](#) name secondary name [bfd-liveness-detection \(LSP\)](#)]

Release Information

Statement introduced in Junos OS Release 19.1R1.

Description

Configure seamless BFD (S-BFD) parameters in the source routing path.

Options

remote-discriminator—Remote discriminator of reflector

Range: 1 through 4294967295

Required Privilege Level

routing

RELATED DOCUMENTATION

[Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution](#) | [676](#)

[bfd-liveness-detection \(LSP\)](#) | [2571](#)

secondary (Protocols MPLS)

Syntax

```
secondary path-name {
  adaptive;
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  bandwidth 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);
  retry-limit number;
  retry-timer seconds;
  select (manual | unconditional);
  standby;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

You can specify secondary paths even if you have not specified any primary paths.

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).

Options

path-name—Name of a path that you created with the **path** statement.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Primary and Secondary LSPs](#) | 565

secondary

Syntax

```
secondary path-name {  
    standby;  
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Specify one or more secondary paths to use for the label switched path (LSP). You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.

Options

path-name —Name of a secondary path that you created with the **path** statement.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Path Protection in an MPLS Network \(CLI Procedure\)](#) | 145

segment

Syntax

```
segment {
  (pop | swap swap);
  description description;
  next-hop next-hop;
  sid-label;
}
```

Hierarchy Level

```
[edit logical-systems name protocols mpls static-label-switched-path],
[edit logical-systems name routing-instances name protocols mpls static-label-switched-path],
[edit protocols mpls static-label-switched-path],
[edit routing-instances name protocols mpls static-label-switched-path]
```

Release Information

Statement introduced in Junos OS Release OS 18.1R1 for MX Series, PTX Series, and QFX Series.

Description

Static segment for segment routing. A static segment is identified by a unique name. This segment type is assigned a segment identifier (SID) which falls under a default range of 100000 through 1048575. The segment has label operation such as pop-and-forward for adjacency segment and swap-and-forward for prefix or node segment. For both types of label operation, the segment is assigned a next hop that specifies the remote IP address if the outgoing interface is a multi-access interface, or the name of the outgoing interface if the interface is a point-to-point interface. Static segment configuration is used to statically configure or provision the adjacency SIDs, node SIDs, and prefix SIDs on transit routers.

Options

pop—Pop the SID label

swap—Swap the SID label to this label

description—Text description of label-switched path

next-hop—IPv4 address or interface of next-hop router

sid-label—Segment identifier (SID) label

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing

RELATED DOCUMENTATION

Static Segment Routing Label Switched Path 632
segment-list 2770
source-routing-path 2782

segment-list

Syntax

```
segment-list name {
  hop-name {
    (loose | strict);
    ip_address IP address;
    label number ;
    label-type node;
  }
  auto-translate {
    protected mandatory;
    unprotected mandatory;
  }
  dynamic;
  compute;
  inherit-label-nexthops;
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing],
[edit protocols source-packet-routing]
```

Release Information

Statement introduced in Junos OS Release 17.4R1 for MX Series and PTX Series with FPC-PTX-P1-A.

ip-address statement introduced in Junos OS Release 18.1R1 on MX Series routers.

inherit-label-nexthops, **node-type**, and **auto-translate** statements introduced in Junos OS Release 19.1R1 on MX Series routers.

dynamic statement introduced in Junos OS Release 19.2R1 on all platforms.

compute, **loose**, and **strict** statements introduced in Junos OS Release 19.1R1-S1 on MX Series routers.


Description

Specify an name to identify the segment routing list (used in traffic engineering policy) and the explicit path for source routing label switched path (LSPs) to traverse through traffic engineering segments. The segment list is essentially a stack of segment identifiers.

Starting in Junos OS release 19.1R1 for MX and PTX Series routers, you can enable a translation service to translate next-hop IP addresses into the corresponding segment identifier (SID) labels. The translation service keeps track of the node reached at each hop.

When configured, the **segment-list** of a segment routing traffic engineering (SR-TE) LSP accepts IP addresses for all the hops along the path. These IP addresses can be either the loopback address of a node, or the

IP address of a link, as identified by the **node-type**. When **auto-translation** is enabled, next hop IP addresses are automatically translated to corresponding SIDs using the translation service. A retry rate can be set for the retry timer at the **source-packet-routing** hierarchy level.



NOTE: The segment list enables BGP and static segment routing LSP to steer traffic based on segment routing policies. When a segment list is used by the protocol BGP, the BGP protocol validates these segment identifiers and selects valid segments for traffic engineering.

Options

<hop-name>—Indicates the next hop in the segment routing traffic engineering policy (SR-TE).

- **ip-address**—Specify the IP address of the hop. For a segment-list to be used by a non-colored segment routing LSP, the first hop must specify an IP address.
- **label**—Specify the SID label of the hop in a segment routing traffic engineering segment list. In static segment routing LSPs, the source routing path uses the segment list only if the second to Nth hop specifies segment identifiers (SID) labels.

NOTE: The range is from 0 to 1,048,576 and is applies to BGP and static segment routing LSPs.

- **label-type**—Use with the option below to indicate that the specified address is the IP address of the node, for example, its loopback address, as opposed to that of a link.
- **node**—Hops that have been specified as **node** are translated to a prefix SID, which can be either a node SID or an anycast SID depending on the type of hop IP address. IP addresses not identified as **node** are consider to be a link.

NOTE: If the first hop is a **node**, for LSP resolution to work correctly, **inherit-label-nexthops** must be enabled at either **source-packet-routing** hierarchy level, or at the relevant **segment-list** hierarchy level.

- **loose | strict**—IP hops specified using router IDs in the sequence can be strict or loose hops. A strict hop must be directly connected to the previous node in the sequence. A loose hop is not necessarily directly connected to the previous node.

NOTE: You can specify only router IDs as loose or strict hop constraints. Labels and other IP addresses are not supported as loose or strict hop constraints in Junos OS Release 19.2R1-S1.

auto-translate— This option must be enabled before a given segment list can use IP addresses instead of SIDs for any hop other than the first hop. In addition, all hops in the segment list must have IP addresses. If any hops on the list have both an IP address and a label configured, the label will be retained. Link addresses are only translated into labels if the preceding node advertises an adjacency SID for the address (otherwise translation fails).

NOTE: In Junos OS Release 19.1R1, for auto-translate to work for OSPF, RSVP for segment routing must be enabled on all participating interfaces.

- **protected**—(Optional) Enable this option to ensure the adjacency SID is eligible to have a backup path, and that a B-flag is set in adjacency SID advertisements. Note that unless **mandatory** is also selected, the choice succeeds regardless.
- **mandatory**—(Optional) Enable this option to have translation fail if any *unprotected* links are found in the hop-list.
- **unprotected**—(Optional) Enable this option to ensure that no backup path is calculated for a specific adjacency SID, and that a B-flag is not set in adjacency SID advertisements. Note that unless **mandatory** is also selected, the choice succeeds regardless.
- **mandatory**—(Optional) Enable this option to have translation fail if any *protected* links are found in the hop-list.

compute—(Optional) Enable use of explicit paths specified in segment list for path computation.

inherit-label-nexthops—Inherit label next hops for first hop in this segment list that have both IP address and label configured in the first hop.

You can configure the **inherit-label-nexthops** statement globally or individually for each segment list.

The **inherit-label-nexthops** statement takes effect only when the segment list first hop has both IP address and SID label present.

If the **inherit-label-nexthops** is not configured at the **[edit protocols source-packet-routing segment-list]** hierarchy, and the first hop in the segment list has both IP address and label specified, the default behavior is to use the IP address.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Static Adjacency Segment Identifier for ISIS</i>
Static Segment Routing Label Switched Path 632
<i>Segment Routing Traffic Engineering at BGP Ingress Peer Overview</i>
Understanding Static Segment Routing LSP in MPLS Networks 632
show spring-traffic-engineering 3578
<i>extended-nexthop-color</i>
source-routing-path 2782
<i>sr-preference-override</i>

select

Syntax

```
select (manual | unconditional);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the conditions under which the path is selected to carry traffic. The **manual** and **unconditional** options are mutually exclusive.

Options

manual—The path is selected for carrying traffic if it is up and stable for at least the revert timer window (potentially before the revert timer has elapsed). Traffic is sent to other working paths if the current path is down or degraded (receiving errors).

unconditional—The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Specifying the Conditions for Path Selection](#) | 567

signal-bandwidth

Syntax

```
signal-bandwidth type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-attributes],  
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the bandwidth encoding of the signal used for path computation and admission control.

Options

type—Configure the type of bandwidth encoding used on the LSP. It can be any of the following values: **10gigether**, **ds1**, **ds3**, **e1**, **e3**, **ethernet**, **fastether**, **gigether**, **stm-1**, **stm-4**, **stm-16**, **stm-64**, **stm-256**, **sts-1**, **vt1-5**, or **vt2**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Signal Bandwidth Type](#) | 1374

signaling

Syntax

```
signaling;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp family inet-mdt],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp family inet-mvpn],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name
  family inet-mdt],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name
  family inet-mvpn],
[edit routing-instances routing-instance-name protocols bgp family inet-mdt],
[edit routing-instances routing-instance-name protocols bgp family inet-mvpn],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet-mdt],
[edit routing-instances routing-instance-name protocols bgp group group-name family inet-mvpn]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Enable signaling in BGP. For multicast distribution tree (MDT) subaddress family identifier (SAFI) NLRI signaling, configure signaling under the **inet-mdt** family. For multiprotocol BGP (MBGP) intra-AS NLRI signaling, configure signaling under the **inet-mvpn** family.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs

site (Layer 2 Circuits)

Syntax

```
site site-name {
    hot-standby;
    site-identifier identifier;
    site-preference preference-value {
        backup;
        primary;
    }
    interface interface-name {
        description text;
        remote-site-id remote-site-ID;
    }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn],
[edit routing-instances routing-instance-name protocols l2vpn]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

hot-standby option introduced in Junos OS Release 14.2 for MX Series routers.

Description

Specify the site name, site identifier, and interfaces connecting to the site. Allows you to configure a remote site ID for remote sites.

Options

hot-standby—Turn on the protector behavior for the site. This keeps backup pseudowire in continuous standby mode and ready for traffic forwarding.

site-identifier *identifier*—Numerical identifier for the site used as a default reference for the remote site ID.

site-name—Name of the site.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Site](#)

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

site-identifier (Layer 2 Circuits)

Syntax

```
site-identifier identifier;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn site site-name],
[edit routing-instances routing-instance-name protocols l2vpn site site-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description

Specify the numerical identifier for the local Layer 2 VPN site.

Options

identifier—The numerical identifier for the Layer 2 VPN site, which can be any number from 1 through 65,534.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Site](#)

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

smart-optimize-timer

Syntax

```
smart-optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

If you want to disable the smart optimizer, you can set it to zero. The **smart-optimize-timer** value in seconds indicates the time before which the LSP is switched back to its primary path in case the primary path becomes available. Otherwise, the time to wait is controlled by the **optimize-timer**, which is usually set to a high value. Some ISPs have the **optimize-timer** set to once a day. Sometimes after the smart optimizer causes the LSP to be placed back on its primary path, the primary path goes down again within 60 minutes. When this happens, the **smart-optimize-timer** is disabled automatically, and the **optimize-timer** (regular path optimization) goes into effect. This is to protect against a flapping link being used.

Default

The smart optimization timer is enabled by default.

Options

seconds—(Optional) Specify the number of seconds to wait before switching an LSP back to its original path. If you do not specify the number of seconds, the default value is used.

Range: 0 through 65,535 seconds

Default: 180 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Smart Optimize Timer for LSPs | 545](#)

[Optimizing Signaled LSPs | 541](#)

[optimize-aggressive | 2714](#)

[optimize-timer | 2717](#)

soft-preemption (Protocols MPLS)

Syntax

```
soft-preemption;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Attempt to establish a new path for a preempted LSP before tearing it down.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS Soft Preemption | 530](#)

source-routing-path

Syntax

```
source-routing-path name {
  binding-sid binding-sid;
  color color;
  metric value;
  no-ingress;
  preference preference;
  primary name {
    weight weight;
  }
  secondary name {
    weight weight;
  }
  sr-preference sr-preference;
  to to;
}
```

Hierarchy Level

```
[edit logical-systems name protocols source-packet-routing],
[edit protocols source-packet-routing]
```

Release Information

Statement introduced in Junos OS Release 17.4R1 for MX Series and PTX Series with FPC-PTX-P1-A. The **metric**, **no-ingress**, and **secondary** statements are introduced in Junos OS Release 18.1R1 for MX Series.

Description

Configure a source routing label switched path (LSP) for steering traffic at an ingress router. Specify a binding segment identifier from the static label range. Configure other parameters such as color, weight, preference, and segment routing (SR) preference for traffic engineering.

Starting with Junos OS Release 18.1R1, static non-colored SR label switched paths (LSPs) for protocol SPRING-TE in an MPLS network. Configure parameters such as destination address, binding SIDs, primary segment, secondary segment, metric, and preference. These SR LSPs do not have a color associated with them. If an ingress route is not required for a non-colored SR LSP then the ingress route installation in inet.3 table can be disabled.

Options

name—Specify a name to identify a source routing path.

binding-sid—Specify the binding label to enable transit functionality for this tunnel. For a non-colored static SR LSP, the binding SID label of protocol SPRING-TE have , by default, a preference of 8 and a metric of 1.

NOTE: This is optional for MPLS networks.

Range: 16 through 1,048,576

color—Specify a color identifier for the tunnel end point.

NOTE: This is only for colored SR LSPs. For non-colored SR LSPs, you do not have to configure the color parameter.

metric—Specify metric for routes downloaded for the non-colored static SR tunnel.

Default:

NOTE: This is the default label range for static LSPs in MPLS networks. You can configure the label range at **[edit protocols mpls label-range static label-range]** hierarchy level.

1000000 through 1048575

Range:

NOTE: This range is for protocol BGP.

1 through 16777215

no-ingress—Disable ingress route that is not required for the non-colored static SR tunnel

preference—Specify the preference for routes downloaded for this tunnel.

primary—Specify a primary segment list for the configured source routing path.

The non-colored static SR LSP can have a maximum of 8 primary paths. incase of multiple operational primary paths, the PFE distributes the traffic over the paths based on the weight configured on the paths. If none of the paths have weights configured then the weights default to 1 making it an ECMP path. the paths become weighted ECMP if the at least one of the paths have a non-zero weight. In both cases , when one or some of the paths fail, the PFE automatically re-balances the traffic over the remaining paths resulting in path protection.

weight *weight_value*—Specify a percentage of the bandwidth with respect to the sum of weights of all paths for the primary segment list. If forwarding interfaces are also configured with weighted ECMP, then Junos OS applies hierarchical weighted ECMP. If the weight percentage is not configured, then only IGP weights are applied on the forwarding interfaces.

secondary—Specify a secondary segment list for the configured non-colored static SR LSP.

sr-preference—Configure a preference for segment routing routes for traffic engineering. BGP chooses a higher preference over a lower preference value.

Range: 0 through 4,294,967,295

to—Specify the IP address of the tunnel end-point

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details..

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>extended-nexthop-color</i>
segment-list 2770
<i>source-packet-routing</i>
<i>sr-preference-override</i>
<i>Segment Routing Traffic Engineering at BGP Ingress Peer Overview</i>

source-routing-path-template

Syntax

```
source-routing-path-template name {
  bfd-liveness-detection;
  ldp-tunneling;
  metric metric;
  no-ingress;
  primary name {
    weight weight;
  }
  secondary name;
  sr-preference sr-preference;
}
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Release Information

Statement introduced in Junos OS Release 19.2R1 on all platforms.

bfd-liveness-detection and **ldp-tunneling** options introduced in Junos OS Release 19.4R1 on all platforms.

Description

Configure a source-routing-path-template for dynamic creation of segment routing label-switched paths (LSPs).

Options

name—Name of the source routing path.

ldp-tunneling—Allow LDP to use this LSP for tunneling. This configuration can be applied to PCE-initiated segment routing LSPs only.

metric—Metric for routes downloaded for this tunnel.

Range: 1 through 16777215

no-ingress—Disable ingress functionality for this tunnel.

primary *name*—Configure a named identifier for the segment-list that describes the primary segment routing path along which the packet is to be routed. This segment list must have the **dynamic** statement enabled for dynamic creation of segment routing LSPs.

weight *weight*—(Optional) Specify the balance factor for this segment list in SR-TE tunnel.

If the forwarding interfaces have weights assigned by IGP, then hierarchical weighted ECMP is applied. When weight is not specified, only the IGP weights are applied on the forwarding interfaces.

secondary name—Configure a named identifier for the segment-list that describes the secondary segment routing path along which the packet is to be routed. This segment list must have the **dynamic** statement enabled for dynamic creation of segment routing LSPs.

sr-preference—Segment routing preference for the segment routing traffic engineered (SPRING-TE) routes. A greater value indicates higher preference.

The preference value of the routes programmed for the segment routing LSP is inherited from the **preference value** statement at the **[edit protocols source-packet-routing]** hierarchy level. When this value is not configured, the default preference value of 8 is used.

Range: 0 through 4294967295

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing

RELATED DOCUMENTATION

[Understanding Static Segment Routing LSP in MPLS Networks](#) | 632

[source-routing-path-template-map](#) | 2787

source-routing-path-template-map

Syntax

```
source-routing-path-template-map {  
  policy [ policy ... ];  
}
```

Hierarchy Level

```
[edit protocols source-packet-routing]
```

Release Information

Statement introduced in Junos OS Release 19.4R1 on all platforms.

Description

Configure a source routing path template map to list the policy statements against which the PCE-initiated LSP should be checked.

Options

policy —Configure policy.

Required Privilege Level

routing

RELATED DOCUMENTATION

[Understanding Static Segment Routing LSP in MPLS Networks](#) | 632

[source-routing-path-template](#) | 2785

splitting-merging

Syntax

```
splitting-merging {
  maximum-member-lsps number;
  maximum-signaling-bandwidth bps;
  merging-bandwidth bps;
  minimum-member-lsps number;
  minimum-signaling-bandwidth bps;
  normalization;
  sampling;
  splitting-bandwidth bps;
  splitting-merging-threshold percent;
}
```

Hierarchy Level

```
[edit protocols mpls container-label-switched-path lsp-name]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Series switches in Junos OS Release 15.1X53-D30.

Description

Perform splitting and merging.

Options

maximum-member-lsps *number*—Number of label-switched paths (LSPs) that a container LSP can have as member LSPs at maximum.

Default: 1

maximum-signaling-bandwidth *bandwidth*—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at maximum after normalization. When **maximum-signaling-bandwidth** is not configured, the value is derived from the **splitting-bandwidth**.

When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the **splitting-bandwidth**.

Default: 1 bps

merging-bandwidth *bandwidth*—Amount of bandwidth in bits per second (bps) that is used for merging during normalization.

Default: 1 bps

minimum-member-lsps *number*—Number of LSPd that a container LSP can have as member LSPs at minimum.

Default: 64

minimum-signaling-bandwidth *bandwidth*—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at minimum after normalization. When **minimum-signaling-bandwidth** is not configured, the value is derived from the **merging-bandwidth**.

When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the **merging-bandwidth**.

Default: 1 bps

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. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [container-label-switched-path](#) | 2579

spring-te (Dynamic Tunnels)

Syntax

```
spring-te {
  destination-networks name;
  source-routing-path-template name {
    (color [ color ... ] | color-any);
  }
}
```

Hierarchy Level

```
[edit routing-options dynamic-tunnels],
```

Release Information

Statement introduced in Junos OS Release 19.2R1 on all platforms.

Description

Enable next hop-base dynamic tunnel mode.

Options

source-routing-path-template *template-name*—Configure template color mapping for segment routing traffic-engineered (SPRING-TE) dynamic LSP parameters.

color—Specify set of color list to be mapped to corresponding SPRING-TE template.

- When enabled, all templates should have color objects defined.
- All destinations are assumed to be colored as well.
- A color can be mapped to only one template at a given point in time.
- Both the **color** and **color-any** statements can coexist. When the two statements are enabled together, preference is given to a specific color template than color-any.
- A colored and non-colored destination cannot co-exist in the same SR-TE configuration.

color-any—Enables mapping of any color to corresponding SPRING-TE template. Only one color-any template can be configured for one SR-TE LSP.

Required Privilege Level

routing

RELATED DOCUMENTATION

Understanding Static Segment Routing LSP in MPLS Networks | 632

srgb-label-range

Syntax

```
srgb-label-range <range-start> <range-end>
```

Hierarchy Level

```
[edit protocols mpls label-range],  
[edit protocols ospf source-packet-routing srgb]
```

Release Information

Statement introduced in Junos OS Release 19.1 for MX Series routers.

Description

SRGB configured under **mpls label-range** is termed as global SRGB. The MPLS label range is based on the start range and the end range. The value of the start range indicates the start of the label range, and the value of the end range along with the value of the start range indicate the end of the label range.

Options

<range-start> <range-end>—Start range and end range of the global SRGB label block.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

source-packet-routing

srlg

Syntax

```
srlg {  
  srlg-name {  
    srlg-cost srlg-cost;  
    srlg-value srlg-value;  
  }  
}
```

Hierarchy Level

```
[edit routing-options],  
[edit logical-systems logical-system-name routing-options]  
[edit protocols mpls interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure Shared Risk Link Group (SRLG) parameters.

Options

srlg-cost *srlg-cost*—Specify a cost for the SRLG ranging from **1** through **65535**.

srlg-value *srlg-value*—Specify a Group ID for the SRLG ranging from **1** through **4294967295**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring SRLG](#) | 278

srlg-cost

Syntax

```
srlg-cost srlg-cost;
```

Hierarchy Level

```
[edit routing-options srlg],  
[edit logical-systems logical-system-name routing-options srlg]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify a cost for the Shared Risk Link Group (SRLG) ranging from **1** through **65535**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Example: Configuring SRLG](#) | **278**

srlg-value

Syntax

```
srlg-value srlg-value;
```

Hierarchy Level

```
[edit routing-options srlg],  
[edit logical-systems logical-system-name routing-options srlg]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify a Group ID for the Shared Risk Link Group (SRLG) ranging from **1** through **4294967295**.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring SRLG](#) | [278](#)

standby

Syntax

```
standby;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |  
    secondary) path-name],  
[edit protocols mpls],  
[edit protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Enable the path to remain up at all times to provide instant switchover if connectivity problems occur.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Hot Standby of Secondary Paths for LSPs | 568](#)

[Configuring Path Protection in an MPLS Network \(CLI Procedure\) | 145](#)

standby

Syntax

```
standby;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |  
    secondary) path-name],  
[edit protocols mpls],  
[edit protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Have the path remain up at all times to provide instant switchover if connectivity problems occur.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Hot Standby of Secondary Paths for LSPs](#) | 568

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.

Statement introduced in 19.4R1 for cRPD instances.

Description

Configure a static LSP.

Options

lsp-name—Name of the path.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Static LSPs](#) | 603

statistics (Protocols MPLS)

Syntax

```
statistics {
  auto-bandwidth (MPLS Statistics);
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-transit-statistics;
  traffic-class-statistics;
  transit-statistics-polling;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

traffic-class-statistics option introduced in Junos OS Release 14.2.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Enable MPLS statistics collection and reporting.

Options

file *filename*—(Optional) Name of the file to receive the output. We recommend that you place MPLS tracing output in the file `mpls-stat` in the `/var/log` directory.

files *number*—(Optional) Maximum number of trace files. When a trace file named *file* reaches its maximum size, it is renamed *file.0*, then *file.1*, and so on, until the maximum number of files is reached. Then, the oldest file is overwritten.

Range: 2 or more

Default: 2 files

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

interval *seconds*—Interval at which to periodically collect statistics.

Range: 1 through 65,535

Default: 300 seconds

no-world-readable—(Optional) Prevent users from reading the log file.

size *size*—(Optional) Maximum size of each file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a file named *file* reaches this size, it is renamed *file.0*. When the *file* again reaches its maximum size, *file.0* is renamed *file.1* and *file* is renamed *file.0*. This renaming scheme continues until the maximum number of files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of files with the **files** option.

world-readable—(Optional) Enable users to read the log file.

Syntax: 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

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. See [CLI Explorer](#).

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 | 242](#)

[Configuring Automatic Bandwidth Allocation for LSPs | 547](#)

swap

Syntax

```
swap out-label;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],  
[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Remove the label at the top of the label stack and replace it with the specified label. Manually assigned incoming labels can have values from 1,000,000 through 1,048,575. This statement is used to configure static LSPs at transit routers.

Options

out-label—Manually assigned outgoing label value.

Range: 0 through 1,048,575

Default: If you do not define the **out-label** option, the original label value remains unchanged.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pop | 2739](#)

[push | 2747](#)

[Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs | 607](#)

switch-away-lsps

Syntax

```
switch-away-lsps;
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols mpls interface interface-name],  
[edit protocols mpls interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 10.2.

Description

(MX Series routers only) Enable you to switch an LSP away from a network node using a bypass LSP. This feature could be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic. Configure this statement only after you have configured and committed the **always-mark-connection-protection-tlv** statement.

The **always-mark-connection-protection-tlv** statement marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. When you configure the **switch-away-lsps** statement, traffic is switched to the bypass LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Switching LSPs Away from a Network Node](#) | 846

switching-type

Syntax

```
switching-type (fiber | lambda | psc-1 | tdm);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name lsp-attributes],  
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the switching method for the LSP. The switching method can be one of the following values:

- **fiber**—Fiber switching
- **lambda**—Lambda switching
- **psc-1**—Packet switching
- **tdm**—Time-division multiplexing (TDM) switching

Default

psc-1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS LSPs for GMPLS](#) | 1372

sync-active-path-bandwidth

Syntax

```
sync-active-path-bandwidth;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name auto-bandwidth],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]
```

Release Information

Statement introduced in Junos OS Release 13.2.

Description

When you have a primary and a secondary path configuration, specify that a path needs to be signaled with the active-path bandwidth when the auto-bandwidth adjustment happens and that the secondary path synchronizes the bandwidth reservations to that of the primary path.

When a primary path fails, bandwidth reservations are made by the secondary path on the links that it uses. If you include the **sync-active-path-bandwidth** statement, the secondary path releases the bandwidth it has reserved and adjusts its bandwidth after the primary path begins carrying traffic.

For example, suppose the active path is a secondary path with a reserved bandwidth of 10 GB as a result of the automatic bandwidth adjustment. Then suppose there is a switchover from the secondary path to the primary path. After some time the primary path reserves 5 GB as a result of a new automatic adjustment. Without the **sync-active-path-bandwidth** statement, the secondary path does not release the 10 GB after a switchover occurs. That bandwidth is wasted. If the **sync-active-path-bandwidth** is included in the configuration, the secondary path adjusts its bandwidth to 5 GB along with the primary path.

Default

When you have a primary and a secondary path configuration, and the primary path fails, bandwidth reservations are made by the secondary path on the links that it uses. When the primary path comes back and the traffic switches over, the secondary path does not release its bandwidth reservations.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling Constrained-Path LSP Computation | 489](#)

[Configuring Explicit-Path LSPs | 686](#)

te-class-matrix

Syntax

```
te-class-matrix {
  tnumber {
    priority priority;
    traffic-class {
      ctnumber priority priority;
    }
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls diffserv-te],
[edit protocols mpls diffserv-te]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.

Default

The default traffic engineering class matrix is:

```
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;
}
```

If you define any of the traffic engineering classes, all the default values are dropped.

Options

ctnumber—Specify the number of the class type. It can be one of four values: **ct0**, **ct1**, **ct2**, or **ct3**.

priority *priority*—Specify the priority of the class type. It can be one of eight values from 0 through 7.

tenumber—Specify the number of the traffic engineering class. It can be one of eight values: **te0**, **te1**, **te2**, **te3**, **te4**, **te5**, **te6**, or **te7**. You must configure the traffic engineering classes in order, starting with **te0**.

traffic-class—Specify the traffic class for the traffic engineering class.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Traffic Engineering Classes](#) | 1202

to

Syntax

```
to address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name bypass],
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls static-label-switched-path lsp-name bypass],
[edit protocols mpls static-label-switched-path lsp-name ingress]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Support for IPv6 addresses in static LSP configurations are provided in Junos OS Release 17.2R1.

Description

Specify the egress router of a dynamic LSP.

Options

address—IPv4 or IPv6 address of the egress router.

NOTE: IPv6 static LSPs are not supported at the `[edit protocols mpls static-label-switched-path lsp-name ingress]` hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Egress Router Address for LSPs](#) | 514

traceoptions (Protocols MPLS)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

ted-export option introduced in Junos OS Release 14.2.

ted-import option introduced in Junos OS Release 14.2.

lsp-history option added in Junos OS Release 15.1.

Description

Configure MPLS tracing options at the protocol level or for a label-switched path.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default MPLS protocol-level tracing options are inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

filename—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

MPLS Tracing Flags

- **all**—Trace all operations
- **autobw-state**—Automatic bandwidth events.
- **connection**—All circuit cross-connect (CCC) activity
- **connection-detail**—Detailed CCC activity
- **cspf**—CSPF computations
- **cspf-link**—Links visited during CSPF computations
- **cspf-node**—Nodes visited during CSPF computations
- **error**—MPLS error packets
- **graceful-restart**—Trace MPLS graceful restart events
- **lsp-history**—Trace LSP history events
- **lsping**—Trace lsping packets and return codes
- **nsr-synchronization**—Trace NSR synchronization events
- **nsr-synchronization-detail**—Trace NSR synchronization events in detail
- **state**—All LSP state transitions
- **static**—Trace static label-switched path
- **ted-export**—Trace leaking of entries from **lsdist.0** table into the traffic engineering database
- **ted-import**—Trace leaking traffic engineering database entries into the **lsdist.0** table
- **timer**—Timer usage

no-world-readable—(Optional) Allow only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Tracing MPLS and LSP Packets and Operations](#) | 1239

traffic-class (delay)

Syntax

```
traffic-class tc-value {
  average-sample-size sample size;
  padding-size size;
  query-interval milliseconds;
  rtt-delay-threshold rtt threshold value;
  twcd-delay-threshold twcd threshold value;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier delay],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier delay],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier delay],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier delay]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure traffic class specific options.

Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the *tc-all|tc-0|tc-1|tc-2|tc-3|tc-4|tc-5|tc-6|tc-7|tc-none* traffic-class values. For each traffic class, you can configure the respective parameters.

To enable traffic-class parameters, configure the **traffic-class-statistics** configuration statement under the **[edit protocol mpls statistic]** hierarchy level.

Options

average-sample-size *sample size*—(Optional) Specify the number of samples used for calculating the average of various metrics.

Default: 5

Range: 1 through 30

padding-size *size*—(Optional) Specify the delay-measurement message length, which is used to calculate the delay experienced by messages of different sizes.

Default: 0

Range: 1 through 1500

query-interval *milliseconds*—Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.

Default: 10 seconds

Range: 1000 through 4294967295 milliseconds

rtt-delay-threshold *rtt threshold value*—Specify the round-trip delay threshold value.

Range: 1 through 4294967295 microseconds

twcd-delay-threshold *twcd threshold value*—Specify the two-way channel delay threshold value.

Range: 1 through 4294967295 microseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements | 273](#)

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

traffic-class (loss)

Syntax

```
traffic-class tc-value {
  average-sample-size sample size;
  loss-threshold loss threshold value;
  loss-threshold-window number of samples for loss threshold;
  measurement-quantity bytes|packets;
  query-interval milliseconds;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier loss],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier loss],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier loss],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier loss]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure traffic class specific options.

Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the *tc-all|tc-0|tc-1|tc-2|tc-3|tc-4|tc-5|tc-6|tc-7|tc-none* traffic-class values. For each traffic class, you can configure the respective parameters.

To enable traffic-class parameters, configure the **traffic-class-statistics** configuration statement under the **[edit protocol mpls statistic]** hierarchy level.

Options

average-sample-size *sample size*—(Optional) Specify the number of samples used for calculating the average of various metrics.

Default: 5

Range: 1 through 30

loss-threshold *loss threshold value*—Specify the threshold value that will be used with **loss-threshold-window** to calculate the loss within specified window size.

Range: 1 through 4294967295

loss-threshold-window *number of samples for loss threshold*—Specify the number of samples used for loss threshold calculation.

Range: 1 through 30

measurement-quantity *bytes|packets*—(Optional) Specify whether packet or byte loss is being measured at the querier.

Default: packets

query-interval *milliseconds*—Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.

Default: 10 seconds

Range: 1000 through 4294967295 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements | 273](#)

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

traffic-class (loss-delay)

Syntax

```
traffic-class tc-value {
  average-sample-size sample size;
  loss-threshold loss threshold value;
  loss-threshold-window number of samples for loss threshold;
  measurement-quantity bytes|packets;
  padding-size size;
  query-interval milliseconds;
  rtt-delay-threshold rtt threshold value;
  twcd-delay-threshold twcd threshold value;
}
```

Hierarchy Level

```
[edit protocols mpls oam performance-monitoring querier loss-delay],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier loss-delay],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier loss-delay],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier loss-delay]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Configure traffic class specific options.

Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the *tc-all|tc-0|tc-1|tc-2|tc-3|tc-4|tc-5|tc-6|tc-7|tc-none* traffic-class values. For each traffic class, you can configure the respective parameters.

To enable traffic-class parameters, configure the **traffic-class-statistics** configuration statement under the **[edit protocol mpls statistic]** hierarchy level.

Options

average-sample-size *sample size*—(Optional) Specify the number of samples used for calculating the average of various metrics.

Default: 5

Range: 1 through 30

loss-threshold *loss threshold value*—Specify the threshold value that will be used with **loss-threshold-window** to calculate loss within specified window size.

Range: 1 through 4294967295

loss-threshold-window *number of samples for loss threshold*—Specify the number of samples used for loss threshold calculation.

Range: 1 through 30

measurement-quantity *bytes|packets*—(Optional) Specify whether packet or byte loss is being measured at the querier.

Default: packets

padding-size *size*—(Optional) Specify the delay-measurement message length, which is used to calculate the delay experienced by messages of different sizes.

Default: 0

Range: 1 through 1500

query-interval *milliseconds*—Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.

Default: 10 seconds

Range: 1000 through 4294967295 milliseconds

rtt-delay-threshold *rtt threshold value*—Specify the round-trip delay threshold value.

Range: 1 through 4294967295 microseconds

twcd-delay-threshold *twcd threshold value*—Specify the two-way channel delay threshold value.

Range: 1 through 4294967295 microseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Pro-Active Loss and Delay Measurements | 273](#)

[On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview | 243](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

traffic-engineering (Protocols MPLS)

Syntax

```
traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],  
[edit protocols mpls]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1 for EX Series switches.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Select whether MPLS performs traffic engineering on BGP destinations only or on both BGP and IGP destinations. Affects only LSPs originating from this routing device, not transit or egress LSPs.

Default

bgp

Options

bgp—On BGP destinations only. Ingress routes are installed in the inet.3 routing table.

bgp-igp—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.

bgp-igp-both-ribs—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.

mpls-forwarding—On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Traffic Engineering for LSPs](#) | 1137

traffic-engineering

Syntax

```
traffic-engineering {
  disable;
}
```

Hierarchy Level

```
[edit protocols ospf | isis]
```

Release Information

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Enable the traffic engineering features of the specified routing protocol.

Default

Traffic engineering is disabled.

Starting in Junos OS release 15.1, traffic engineering is enabled by default whenever the IS-IS protocol is enabled. You can disable it by including the **disable** statement at the **[edit protocols isis traffic-engineering]** hierarchy level. For the EX3300, EX4200, EX4500, EX4550, EX8200 and XRE200, you can disable traffic engineering starting in Junos OS release 15.1R7.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

[Configuring an OSPF Network \(J-Web Procedure\)](#)

[MPLS Applications User Guide](#)

traffic-engineering (Protocols BGP)

Syntax

```
traffic-engineering {
    unicast;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp family],
[edit logical-systems logical-system-name protocols bgp group group-name family],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name
    family],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name
    neighbor address family],
[edit protocols bgp family],
[edit protocols bgp group group-name family],
[edit protocols bgp group group-name neighbor address family],
[edit routing-instances routing-instance-name protocols bgp family],
[edit routing-instances routing-instance-name protocols bgp group group-name family],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address family]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced on QFX5100 switches in Junos OS Release 15.1

Statement introduced on QFX10000 switches in Junos OS Release 17.1.

Description

Enable traffic engineering address family. This generates a multiprotocol address family indicator (AFI) and a subsequent address family identifier (SAFI) to be negotiated with the BGP peers.

The BGP network layer reachability information (NLRI) information is exchanged between the peers only when the traffic engineering AFI and SAFI are shared between them. If the peers do not agree on the use of the AFI and SAFI, the connection between the peers is terminated.

Options

unicast—Include BGP-TE NLRI.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring Link State Distribution Using BGP](#) | 1163

transit-lsp-association

Syntax

```
transit-lsp-association transit-association-lsp-group-name {
  from-1 address-of-associated-lsp-1;
  from-2 address-of-associated-lsp-2;
  lsp-name-1 name-of-associated-lsp-1;
  lsp-name-2 name-of-associated-lsp-2;
}
```

Hierarchy Level

[edit protocols mpls]

Release Information

Statement introduced in Junos OS Release 12.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Associate two label-switched paths (LSPs) at a transit node to configure a path for sending and receiving GAL and G-Ach messages for MPLS-TP OAM.

Options

transit-association-lsp-group-name—Name of the transit association LSP group.

from-1 address-of-associated-lsp-1—Address of the first associated LSP.

from-2 address-of-associated-lsp-2—Address of the second associated LSP.

lsp-name-1 name-of-associated-lsp-1—Name of the first associated LSP.

lsp-name-2 name-of-associated-lsp-1—Name of the second associated LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring the MPLS Transport Profile for OAM](#) | 1222

ultimate-hop-popping

Syntax

```
ultimate-hop-popping;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path label-switched-path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path label-switched-path-name]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Description

Enable ultimate-hop popping on LSPs. Configure this statement on the device at the LSP ingress. In ultimate-hop popping, the MPLS label is popped from the IP packet at the PE router. The IP address is checked in a second address lookup (also at the PE router), and then the packet is forwarded to its destination.

Be aware of the following platform requirements and restrictions:

- UHP LSPs using VT interfaces—Supported on all M Series, MX Series, T Series, and TX Matrix routers.
- UHP LSPs using LSI interfaces—Supported on MX 3D Series routers only.
- UHP LSP requirements for the egress PE device—For M Series and T Series routers, a VT interface is needed.
- UHP LSPs and Layer 3 VPNs—UHP LSPs are supported for Layer 3 VPNs configured on MX 3D Series routers only.
- UHP LSPs and VPLS—UHP LSPs are supported for VPLS configured on MX 3D Series routers only. You must configure the *no-tunnel-services* statement at the **[edit routing-instances *routing-instance-name* protocols vpls]** hierarchy level.

Default

Ultimate-hop popping is disabled by default on LSPs. Penultimate-hop popping is the default behavior. In penultimate-hop popping, the final MPLS label is popped from the IP packet at the last provider router in the network before being forwarded to the PE router. The PE router receives the packet and checks the IP address, and then the packet is forwarded to its destination.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Ultimate-Hop Popping for LSPs | 599](#)

[explicit-null | 2611](#)

vrf-table-label

Syntax

```
vrf-table-label {  
    source-class-usage;  
    static;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],  
[edit routing-instances routing-instance-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for the **source-class-usage** statement added in Junos OS Release 9.3.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Statement introduced in Junos OS Release 15.1F5 and 16.1R2 for PTX5000 routers with third-generation FPCs installed.

Statement introduced in Junos OS Release 15.1F6 and 16.1R2 for PTX3000 routers with third-generation FPCs.

Statement introduced in Junos OS Release 16.1X65 and 17.2R1 for PTX1000 routers.

Statement introduced in cRPD Release 19.4R1.

Support for the **static** statement added in Junos OS Release 17.2.

Description

Map the inner label of a packet to a specific VPN routing and forwarding (VRF) instance. This allows the examination of the encapsulated IP header. The first lookup is done on the VPN label to determine which VRF instance to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts.

When you include the **vrf-table-label** statement in the configuration of a VRF routing instance, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table. Any routes in the VRF routing table are advertised with the LSI logical interface label allocated for the VRF routing table. When packets destined for the VRF routing instance arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table.

All routes in a VRF routing instance configured with this option are advertised with one label allocated per VRF.

NOTE:

- The **vrf-table-label** statement is supported on PTX5000 and PTX3000 routers only when third-generation FPCs are installed on the router and **enhanced-ip** command is configured on the chassis.
- Starting in Junos OS Release 17.2, you can configure the **enhanced-ip** command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can also separate the MPLS labels used for different label spaces which provides more flexibility and scalability. The **vrf-table-label** space is increased to at least 16,000, if the platform can support the scale.

Options

The remaining statements are explained separately.

Range: 16 through 1,048,575 for static label value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Filtering Packets in Layer 3 VPNs Based on IP Headers

Configuring EXP-Based Traffic Classification for VPLS

Load Balancing and IP Header Filtering for Layer 3 VPNs

RSVP Configuration Statements

IN THIS CHAPTER

- admin-group | 2831
- aggregate (Protocols RSVP) | 2832
- authentication-key (Protocols RSVP) | 2834
- bandwidth (Protocols RSVP) | 2835
- bypass (Signaled LSP) | 2837
- bypass (Static LSP) | 2839
- chained-composite-next-hop | 2840
- class-of-service (Protocols RSVP) | 2843
- destination-networks | 2844
- devices | 2845
- disable (Protocols RSVP) | 2846
- dynamic-bidirectional-transport | 2847
- fast-reroute (Protocols RSVP) | 2848
- graceful-deletion-timeout | 2849
- graceful-restart (Protocols RSVP) | 2850
- hello-acknowledgements | 2851
- hello-interval (Protocols RSVP) | 2852
- hop-limit | 2853
- interface (Protocols RSVP) | 2855
- keep-multiplier | 2857
- label-switched-path-template (Multicast) | 2859
- link-protection (RSVP) | 2861
- load-balance (Protocols RSVP) | 2863
- max-bypasses | 2864
- no-local-reversion | 2865
- node-hello | 2867
- no-adjacency-down-notification (Protocols IS-IS) | 2869
- no-authentication-check (Protocols RSVP) | 2870

- no-cspf (Protocols RSVP) | 2871
- no-interface-hello | 2872
- no-neighbor-down-notification | 2873
- no-node-id-subobject | 2874
- no-p2mp-sublsp | 2875
- no-enhanced-frr-bypass (Protocols RSVP) | 2876
- node-link-protection (Protocols MPLS) | 2877
- optimize-timer (Protocols RSVP) | 2878
- path (Protocols RSVP) | 2879
- peer-interface (Protocols RSVP) | 2880
- pop-and-forward (Protocols RSVP) | 2881
- preemption | 2882
- priority (Protocols RSVP) | 2884
- refresh-time | 2885
- reliable | 2886
- rsvp | 2887
- rsvp-te (Routing Options) | 2888
- setup-protection | 2889
- soft-preemption (Protocols RSVP) | 2890
- static-label-switched-path | 2891
- subscription | 2893
- traceoptions (Protocols RSVP) | 2895
- transit | 2898
- tunnel-services (RSVP) | 2900
- ultimate-hop-popping | 2901
- update-threshold | 2903

admin-group

Syntax

```
admin-group {
  exclude [ group-names ];
  include-all [ group-names ];
  include-any [ group-names ];
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Description

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.

Options

exclude *group-names*—Specify the administrative groups to exclude for a bypass LSP.

include-all *group-names*—Specify the administrative groups whose links the bypass LSP must traverse.

include-any *group-names*—Specify the administrative groups whose links the bypass LSP can traverse.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Administrative Groups for Bypass LSPs](#) | 873

aggregate (Protocols RSVP)

Syntax

```
(aggregate | no-aggregate);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],  
[edit protocols rsvp peer-interface peer-interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Control the use of RSVP aggregate messages on an interface or peer interface, as described below.

NOTE: Starting in Junos OS Release 15.2, the **aggregate** statement is deprecated at the **[edit protocols rsvp interface *interface-name*]** and **[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*]** hierarchy levels, as RSVP message aggregation is enabled by default.

- **aggregate**—Use RSVP aggregate messages.
- **no-aggregate**—Do not use RSVP aggregate messages.

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.

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.

To have refresh reduction and reliable delivery, you must include the **aggregate** and **reliable** statements.

Default

Aggregation is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

RSVP Refresh Reduction 823
Configuring RSVP Refresh Reduction 829
reliable 2886

authentication-key (Protocols RSVP)

Syntax

```
authentication-key key;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],
[edit protocols rsvp],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp peer-interface peer-interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface. To authenticate node hellos or remote messages between the Point of Local Repair (PLR) to the Merge Point (MP), enable **authentication-key** at the **[edit protocols rsvp]** hierarchy level.

RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.

All routers that are connected to the same IP subnet must use the same authentication scheme and password.

Options

key—Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring RSVP Authentication](#) | 832

bandwidth (Protocols RSVP)

Syntax

```
bandwidth bps;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

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.

Default

The hardware raw bandwidth is used.

Options

bps—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations **k** (for a thousand), **m** (for a million), or **g** (for a billion [also called a thousand million]).

Range: Any positive integer

Default: 0 (no bandwidth is reserved)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Bandwidth for Bypass LSPs | 874](#)

[Configuring Link Protection on Interfaces Used by LSPs | 871](#)

[Configuring Bypass LSPs | 872](#)

bypass (Signaled LSP)

Syntax

```
bypass bypass-name {
  bandwidth bps;
  description text;
  hop-limit number;
  no-cspf;
  path address <strict | loose>;
  priority setup-priority reservation-priority;
  subscription subscription-percentage;
  to address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection]
```

Release Information

Statement introduced before Junos OS Release 7.4.

description option was added in Junos OS Release 10.4.

subscription *subscription-percentage* option introduced in Junos OS Release 19.2R1 on all platforms.

Description

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.

If you specify the **bandwidth**, **hop-limit**, and **path** statements for the bypass LSP, these values take precedence over the values configured at the **[edit protocols rsvp interface *interface-name* link-protection]** hierarchy level. The other attributes (**subscription**, **no-node-protection**, and **optimize-timer**) are inherited from the general constraints.

Options

bypass-name—(Required) Specify a name for the bypass LSP. The name can be up to 64 characters.

description—Provides a textual description of the bypass LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show mpls lsp bypass detail** command and has no effect on the operation of the bypass LSP. The description text can be no more than 80 characters in length.

subscription *subscription-percentage*—(Optional) Specify the subscription percentage per manual bypass LSP. The subscription percentage configured under a particular manual bypass LSP overrides the subscription percentage configured commonly for all manual bypass LSPs under an interface.

Range: 0 through 65000

to *address*—(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Bypass LSPs](#) | 872

bypass (Static LSP)

Syntax

```
bypass bypass-name {  
    bandwidth bps;  
    description string;  
    next-hop (address | interface-name | address/interface-name);  
    push out-label;  
    to address;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name],  
[edit protocols mpls static-label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 10.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Static LSPs](#) | 603

chained-composite-next-hop

Syntax

```
chained-composite-next-hop {  
    ingress;  
    transit;  
}
```

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 5G Universal Routing Platforms 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 explicitly configure chained composite next hops, include the **enhanced-ip** statement at the **[edit chassis network-services]** hierarchy level. However, take the following into consideration when enabling the **enhanced-ip** mode:

- Non-service DPCs do not work with enhanced network services mode options. Only MPCs, MS-DPCs, and MS-MPCs provide support for the **enhanced-ip** configuration.
- If you configure chained composite next hops on MX Series routers with both MPCs, and DPCs or DPCEs, the network services mode must be changed from **enhanced-ip** to **ethernet** for the DPC or DPCE to come online.

You can verify the FPC status in such cases, using the **show chassis fpc** command output, where the DPCs and DPCEs are marked as **FPC misconfiguration**.

NOTE:

- When chained composite next hops are enabled on a device, the BGP connections are reset.
- On MX Series routers with DPCs or DPCEs, only the *l3vpn* option is supported under the *ingress* configuration statement; all other configuration options and functionality are not supported.
- The **transit** statement and the associated functionality is supported only on PTX Packet Transport Routers and QFX10000 switches.
- 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
```

- Starting with Junos OS Release 17.2, you cannot configure **chained-composite-next-hop ingress l3vpn extended-space** on a logical system.

The remaining statements are explained separately. See [CLI Explorer](#).

Default

This statement is disabled by default.

Options

ingress—Enable or disable composite chained next hop for ingress traffic.

transit—(PTX and QFX10000) Enable or disable composite chained next hop for transit traffic. 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. See [CLI Explorer](#).

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 | 1426](#)

Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs

ingress

[transit \(Chained Composite Next Hops\) | 3352](#)

class-of-service (Protocols RSVP)

Syntax

```
class-of-service cos-value;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

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.

The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.

Options

cos-value—CoS value. A higher value typically corresponds to a higher level of service.

Range: 0 through 7

Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Class of Service for Bypass LSPs](#) | 875

destination-networks

Syntax

```
destination-networks prefix;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options dynamic-tunnels
  tunnel-name],
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name rsvp-te entry],
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],
[edit routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name],
[edit routing-instances routing-instance-name routing-options dynamic-tunnels tunnel-name rsvp-te entry],
[edit routing-options dynamic-tunnels tunnel-name],
[edit routing-options dynamic-tunnels tunnel-name rsvp-te entry]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created.

Options

prefix—Destination prefix of the network.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring GRE Tunnels for Layer 3 VPNs](#)

[Configuring Dynamic Tunnels](#)

[Configuring RSVP Automatic Mesh](#) | 850

devices

Syntax

```
devices device-names;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specifies one of the virtual tunnel (VT) interfaces to de-encapsulate the egress traffic for ultimate-hop popping on point-to-multipoint LSPs. If no device is specified, the selection process is performed automatically.

Default

The device selection process is performed automatically if no device is configured. Junos OS selects one of the available VT interfaces to de-encapsulate the egress traffic.

Options

device-names—Specify which VT interfaces are used to handle the RSVP traffic.

Range: 0 to 8 devices

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs | 859](#)

Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs

disable (Protocols RSVP)

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit logical-systems logical-system-name protocols rsvp graceful-restart],
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],
[edit protocols rsvp],
[edit protocols rsvp graceful-restart],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp peer-interface peer-interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.

Default

RSVP is enabled on interfaces and peer interfaces configured with the RSVP **interface** statement. RSVP graceful restart is enabled on the router. Link protection is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Minimum RSVP Configuration | 827](#)

[Configuring RSVP Graceful Restart | 883](#)

[Configuring Link Protection on Interfaces Used by LSPs | 871](#)

dynamic-bidirectional-transport

Syntax

```
dynamic-bidirectional-transport {  
    template template;  
}
```

Hierarchy Level

```
[edit protocols rsvp peer-interface peer-interface-name]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Enable the dynamic setup of associated bidirectional packet LSP for transporting non-packet Generalized Multiprotocol Label Switching (GMPLS) label-switched path (LSP).

Options

template *template*—Name of the template for the dynamic bidirectional packet LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

fast-reroute (Protocols RSVP)

Syntax

```
fast-reroute optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement added in Junos OS Release 7.5.

Statement introduced in Junos OS Release 14.1 for the QFX Series.

Description

Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.

Options

seconds—Specify the number of seconds between fast reroute detour LSP optimizations.

Range: 0 through 65,535 seconds

Default: 0 (disabled)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Optimization Interval for Fast Reroute Paths](#) | 483

graceful-deletion-timeout

Syntax

```
graceful-deletion-timeout seconds;
```

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

Specify the time, in seconds, before completing graceful deletion of signaling.

Options

seconds—Time before completing graceful deletion of signaling.

Range: 1 through 300 seconds

Default: 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Graceful Deletion Timeout Interval](#) | 1376

graceful-restart (Protocols RSVP)

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  maximum-helper-recovery-time seconds;
  maximum-helper-restart-time seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options],
[edit protocols rsvp],
[edit routing-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure graceful restart on the router. You must configure the **graceful-restart** statement at the **[edit routing-options]** hierarchy level to enable graceful restart on the router.

Options

disable—Disable graceful restart on the router or for RSVP.

helper-disable—Disable RSVP graceful restart helper mode (this option is only available at the **[edit protocols rsvp]** hierarchy level).

Default: Helper mode is enabled by default.

maximum-helper-recovery-time *seconds*—The maximum length of time the router stores the state of neighboring routers when they undergo a graceful restart. The value applies to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.

Default: 180 seconds

Range: 1 through 3600 seconds

maximum-helper-restart-time *seconds*—The maximum length of time the router waits between when it discovers that a neighboring router has gone down and when it declares the neighbor down. This value is applied to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.

Default: 20 seconds

Range: 1 through 1800 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring RSVP Graceful Restart](#) | 883

hello-acknowledgements

Syntax

```
hello-acknowledgements;
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 10.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once hello acknowledgments are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator.

Default

Hello acknowledgments are disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Hello Acknowledgments for Nonsession RSVP Neighbors](#) | 846

hello-interval (Protocols RSVP)

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],  
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],  
[edit protocols rsvp interface interface-name],  
[edit protocols rsvp peer-interface peer-interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable the sending of hello packets on the interface.

Options

seconds—Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface.

Range: 1 through 60 seconds

Default: 9 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the RSVP Hello Interval](#) | 831

hop-limit

Syntax

```
hop-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name fast-reroute],
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name (primary |
secondary) path-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path lsp-name],
[edit protocols mpls label-switched-path lsp-name fast-reroute],
[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:

- LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.
- Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.
- Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.

Options

number—Maximum number of hops.

Range: 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)

Default: 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Fast Reroute | 479](#)

[Limiting the Number of Hops in LSPs | 546](#)

[Configuring the Hop Limit for Bypass LSPs | 875](#)

interface (Protocols RSVP)

Syntax

```

interface interface-name {
    disable;
    (aggregate | no-aggregate);
    authentication-key key;
    bandwidth bps;
    hello-interval seconds;
    link-protection {
        disable;
        admin-group {
            exclude [ group-names ];
            include-all [ group-names ];
            include-any [ group-names ];
        }
        bandwidth bps;
        bypass bypass-name {
            bandwidth bps {
                ct0 bps;
                ct1 bps;
                ct2 bps;
                ct3 bps;
            }
            description text;
            class-of-service cos-value;
            hop-limit number;
            no-cspf;
            path address <strict | loose>;
            priority setup-priority reservation-priority;
            to address;
        }
        class-of-service cos-value;
        hop-limit number;
        max-bypasses number;
        no-cspf;
        no-node-protection;
        optimize-timer seconds;
        path address <strict | loose>;
        priority setup-priority reservation-priority;
        subscription percentage;
    }
    (reliable | no-reliable);
    subscription percentage {

```



```

    ct0 percentage;
    ct1 percentage;
    ct2 percentage;
    ct3 percentage;
  }
  update-threshold threshold;
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]

```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable RSVP on one or more router interfaces.

Default

RSVP is disabled on all interfaces.

Options

interface-name—Name of an interface. To configure all interfaces, specify **all**. For details about specifying interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Minimum RSVP Configuration](#) | 827

keep-multiplier

Syntax

```
keep-multiplier number;
```

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

Specify a value used by RSVP to calculate timer values for network outages, and declare that a reservation or a neighbor is down. It indicates the number of messages that can be lost before a particular state is declared stale and must be deleted. The keep multiplier directly affects the lifetime of an RSVP state.

Starting in Junos OS Release 16.1, for MX Series routers and PTX Series routers, the default RSVP message refresh time, to which this multiplier is applied, has increased from 30 seconds to 20 minutes. The higher message refresh time provides support for RSVP Refresh Reduction Extensions, and improved scaling for MPLS traffic-engineered LSPs, as defined in RFC 2961. The changes are backward compatible so if any nodes in the two-hop neighborhood do not support the higher refresh time, the updated node will automatically fall back to the previous default refresh time to prevent error or tear down messages.

Options

number—Multiplier value.

Range: 1 through 255

Default: 3

NOTE: For MX Series routers and PTX Series routers (running Junos OS release 16.1 or later), this multiplier is applied to a default refresh time of 20 minutes. In earlier Junos OS releases, and for other platforms, the default refresh time remains 30 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Timers for RSVP Refresh Messages](#) | 851

label-switched-path-template (Multicast)

Syntax

```
label-switched-path-template {
  (default-template | lsp-template-name);
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-tunnel rsvp-te],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-tunnel ingress-replication
  label-switched-path],
[edit logical-systems logical-system-name routing-instances routing-instance-name provider-tunnel selective group
  address source source-address rsvp-te],
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name rsvp-te entry-name],
[edit protocols mvpn inter-region-segmented template template-name region region-name ingress-replication
  label-switched-path],
[edit protocols mvpn inter-region-segmented template template-name region region-name rsvpe-te],
[edit protocols mvpn inter-region-template template template-name all-regions ingress-replication
  label-switched-path],
[edit protocols mvpn inter-region-template template template-name all-regions rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel ingress-replication label-switched-path],
[edit routing-instances routing-instance-name provider-tunnel rsvp-te],
[edit routing-instances routing-instance-name provider-tunnel selective group address source source-address rsvp-te],
[edit routing-options dynamic-tunnels tunnel-name rsvp-te entry-name]
[edit routing-instances instance-name provider-tunnel]
```

Release Information

Statement introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 18.2. under the heirarchy level [edit routing-instances *instance-name* provider-tunnel]

Description

Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs. This feature can be used for a number of applications, including point-to-multipoint LSPs, flooding VPLS traffic, configuring ingress replication for IP multicast using MBGP MVPNs, and to enable RSVP automatic mesh. There is no default setting for the **label-switched-path-template** statement, so you must configure either the default-template using the **default-template** option, or you must specify the name of your preconfigured LSP template.

Options

default-template—Specify that the default LSP template be used for the dynamically generated LSPs.

lsp-template-name—Specify the name of an LSP to be used as a template for the dynamically generated LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs

Configuring Point-to-Multipoint LSPs for an MBGP MVPN

Configuring Dynamic Point-to-Multipoint Flooding LSPs

[Configuring RSVP Automatic Mesh](#) | 850

link-protection (RSVP)

Syntax

```
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;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the **link-protection** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level. You can configure single or multiple bypasses for protected interface.

Default

Link protection is disabled.

Options

no-node-protection—Disable node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately. See [CLI Explorer](#).

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 | 871](#)

[link-protection \(Dynamic LSPs\) | 2662](#)

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](#) | 848

max-bypasses

Syntax

```
max-bypasses number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],  
[edit protocols rsvp interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Range modified in Junos OS Release 9.3.

Description

Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure **max-bypasses**, you must also configure the **bandwidth** statement.

Options

number—Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured.

Range: 0 through 99

Default: 1

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Maximum Number of Bypass LSPs | 876](#)

no-local-reversion

Syntax

```
local-reversion;  
no-local-reversion;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Description

Disable RSVP local revertive mode as specified in RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*.

NOTE: For Junos OS Release 16.1 running on MX Series or PTX Series routers, **no-local-reversion** is enabled by default, that is, local reversion is not running, and the statement has been deprecated. To enable local reversion, use the **local-reversion** statement.

RSVP local revertive mode is supported on all Juniper Networks routers running 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.

The following information can also be found in RFC 4090. Refer to the 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:

- 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.
- Global revertive mode—The ingress router of each tunnel is responsible for reoptimizing the LSPs that used the failed path. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or IS-IS LSPs, and timers. This re-optimization process can proceed as soon as the failure is detected. It is not tied to the restoration of the failed path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

node-hello

Syntax

```
(node-hello | no-node-hello);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable 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 OS uses interface-based RSVP hellos; node-ID based RSVP hellos are disabled. If you have not enabled RSVP node IDs on the router, Junos OS does not accept any node-ID hello packets.

NOTE: For Junos OS Release 16.1 running on MX Series or PTX Series routers, when using enhanced FRR, node-ID based hellos are enabled by default. Disabling the enhanced fast reroute (FRR) profile by using the **no-enhanced-frr-bypass** command also disables the **node-hello** command. To re-enable node hellos after the enhanced FRR profile has been disabled, use the **node-hello** command.

NOTE: If **link-protection** is enabled, remote node hellos that are initiated by the Point of Local Repair (PLR) to Node Protecting Merge Point (NP-MP) are enabled automatically. Similarly, if **no-enhanced-frr-profile** is enabled (that is, enhanced FRR is off), remote node hellos are automatically disabled. There is no command to explicitly enable or disable remote-node hellos.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring RSVP Node-ID Hellos	835
no-enhanced-frr-bypass	2876

no-adjacency-down-notification (Protocols IS-IS)

Syntax

```
no-adjacency-down-notification;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols isis interface interface-name],  
[edit protocols isis interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 8.0.

Description

Disable adjacency down notification for IS-IS to allow for migration from IS-IS to OSPF without disruption of the RSVP neighbors and associated RSVP-signaled label-switched paths (LSPs).

Whenever IS-IS is deactivated, the IS-IS adjacencies are brought down. IS-IS signals to RSVP to bring down any RSVP neighbors associated with the IS-IS adjacencies, and this further causes the associated LSPs signaled by RSVP to go down as well.

A similar process occurs whenever OSPF is deactivated. The OSPF neighbors are brought down. OSPF signals to RSVP to bring down any of the RSVP neighbors associated with the OSPF neighbors, and this further causes the associated LSPs signaled by RSVP to go down as well.

If you need to migrate from IS-IS to OSPF or from OSPF to IS-IS, the internal gateway protocol (IGP) notification to RSVP for an adjacency or neighbor down event needs to be ignored. Using the **no-adjacency-down-notification** or **no-neighbor-down-notification** statements, you can disable IS-IS adjacency down notification or OSPF neighbor down notification, respectively, until the migration is complete. The network administrator is responsible for configuring the statements before the migration, and then removing them from the configuration afterward, so that IGP notification can function normally.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[no-neighbor-down-notification](#) | 2873

no-authentication-check (Protocols RSVP)

Syntax

```
no-authentication-check;
```

Hierarchy Level

```
[edit protocols rsvp],
```

Release Information

Statement introduced in Junos OS Release 18.4R1.

Description

Skip authentication check for received messages.

Required Privilege Level

routing

RELATED DOCUMENTATION

[RSVP Authentication](#) | **816**

no-cspf (Protocols RSVP)

Syntax

```
no-cspf;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],  
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],  
[edit protocols rsvp interface interface-name link-protection],  
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Release Information

Statement introduced in Junos OS Release 7.5.

Description

Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.

Default

CSPF is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling CSPF for Bypass LSPs](#) | 876

no-interface-hello

Syntax

```
no-interface-hello;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Explicitly disable RSVP interface hellos globally on the router.

NOTE: For Junos OS Release 16.1 running on MX Series or PTX Series routers, the behavior of this statement has changed. On these platforms, rather than disabling RSVP interface hellos globally, the **no-interface-hello** command triggers a switch back to the previous profile for all label-switched paths (LSPs.)

This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the [hello-interval](#) statement. This configuration disables RSVP interface hellos globally but enables RSVP interface hellos on the specified interface. This configuration might be necessary in a heterogeneous network where some devices support RSVP node-ID hellos and other devices support RSVP interface hellos.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring RSVP Node-ID Hellos | 835](#)

[hello-interval \(Protocols RSVP\) | 2852](#)

no-neighbor-down-notification

Syntax

```
no-neighbor-down-notification;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id interface interface-name],  
[edit protocols ospf area area-id interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 8.0.

Description

Disable neighbor down notification for OSPF to allow for migration from OSPF to IS-IS without disruption of the RSVP neighbors and associated RSVP-signaled LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

no-node-id-subobject

Syntax

```
no-node-id-subobject;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Disable the record route object (RRO) node-ID subobject for compatibility with earlier versions of Junos OS.

NOTE: For Junos OS Release 16.1 running on MX Series or PTX Series routers, the behavior of this statement has changed. On these platforms, rather than disabling the record route object (RRO) node ID sub-object, the **no-node-id-subobject** command triggers a switch back to the previous profile for all label-switched paths (LSPs).

To interoperate with other vendors' equipment, Junos OS supports the RRO node-ID subobject for use in inter-AS link and node protection configurations.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Inter-AS Node and Link Protection](#) | 880

no-p2mp-sublsp

Syntax

```
no-p2mp-sublsp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 9.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Reject Resv messages that include the S2L_SUB_LSP object. By default, Resv messages that include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both Junos OS Release 9.2 and later and Junos OS Release 9.1 and earlier, it is necessary to configure the **no-p2mp-sublsp** statement on devices running Junos OS Release 9.2 and later to ensure that point-to-multipoint LSPs function properly.

Default

Resv messages that include the S2L_SUB_LSP object are accepted.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases](#) | 728

no-enhanced-frr-bypass (Protocols RSVP)

Syntax

```
no-enhanced-frr-bypass;
```

Hierarchy Level

```
[edit protocols rsvp ]
```

Release Information

Statement introduced in Junos OS Release 15.2R1.

Description

Enable **no-enhanced-frr-bypass** to turn off all Fast reroute (FRR) facility protection enhancements, which includes improved LSP scaling and enhanced RSVP message handling, and reduce the default refresh time to 30 seconds.

FRR is enabled by default for MX Series and PTX Series routers starting in Junos OS Release 15.2R1.

Default

This feature, **no-enhanced-frr-bypass**, is disabled by default. That is, enhanced FRR is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [RSVP Refresh Reduction](#) | 823

node-link-protection (Protocols MPLS)

Syntax

```
node-link-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name],  
[edit protocols mpls label-switched-path lsp-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the **link-protection** statement at the **[edit protocols rsvp interface *interface-name*]** hierarchy level.

Default

Node and link protection is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Node Protection or Link Protection for LSPs | 879](#)

[MPLS Feature Support on QFX Series and EX4600 Switches | 26](#)

[Understanding Interprovider and Carrier-of-Carriers VPNs | 1647](#)

optimize-timer (Protocols RSVP)

Syntax

```
optimize-timer seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],  
[edit protocols rsvp interface interface-name link-protection]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both.

Options

seconds—Specify the number of seconds between optimizations.

Range: 0 through 65,535 seconds

Default: 0 (disabled)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Optimization Interval for Bypass LSPs](#) | 877

path (Protocols RSVP)

Syntax

```
path address <strict | loose>;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],
[edit protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name link-protection bypass bypass-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.

Default

No path is configured. CSPF automatically calculates the path the bypass LSP takes.

Options

address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is **loose**. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.

Default: If you do not specify any routers explicitly, no routing limitations are imposed on the bypass LSP.

loose—(Optional) The next address in the **path** statement is loose. The LSP can traverse other routers before reaching this router.

Default: **strict**

strict—(Optional) The LSP must go to the next address specified in the **path** statement without traversing other nodes. This is the default.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring an Explicit Path for Bypass LSPs](#) | 878

peer-interface (Protocols RSVP)

Syntax

```
peer-interface peer-interface-name {  
    disable;  
    (aggregate | no-aggregate);  
    authentication-key key;  
    dynamic-bidirectional-transport template template;  
    hello-interval seconds;  
    (reliable | no-reliable);  
}
```

Hierarchy Level

```
[edit protocols rsvp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

dynamic-bidirectional-transport template *template* option introduced in Junos OS Release 14.2.

Description

Configure the name of the LMP peer device.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring RSVP and OSPF for LMP Peer Interfaces](#) | 1369

pop-and-forward (Protocols RSVP)

Syntax

```
pop-and-forward {  
    application-label depth depth;  
}
```

Hierarchy Level

```
[edit logical-systems name protocols rsvp],  
[edit logical-systems name routing-instances name protocols rsvp],  
[edit protocols rsvp],  
[edit routing-instances name protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 18.1R1 on MX Series routers, PTX Series routers, and vMX.

Description

Specify RSVP pop-and-forward LSP tunnel-specific global parameters. The application label depth (AppLD) value must be configured uniformly across the RSVP-TE network.

Options

application-label depth *depth*—Specify the maximum number of service labels.

Range: 0 through 3

Default: 1

Required Privilege Level

routing

RELATED DOCUMENTATION

[Pop-and-Forward LSP Configuration](#) | 797

[show rsvp pop-and-forward](#) | 3387

[pop-and-forward \(Protocols MPLS\)](#) | 2740

preemption

Syntax

```
preemption {  
  (aggressive | disabled | normal);  
  soft-preemption {  
    cleanup-timer seconds;  
  }  
}
```

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

Control RSVP session preemption.

Default

normal

Options

aggressive—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.

disabled—Do not preempt RSVP sessions.

normal—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

priority (Protocols RSVP)

Syntax

```
priority setup-priority reservation-priority;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],  
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection bypass bypass-name],  
[edit protocols rsvp interface interface-name link-protection],  
[edit protocols rsvp interface interface-name link-protection bypass bypass-name],
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.

Options

reservation-priority—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.

Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.

Default: 0 (Once the session is set up, no other session can preempt it.)

setup-priority—Setup priority.

Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.

Default: 7 (The session cannot preempt any existing sessions.)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Priority and Preemption for Bypass LSPs | 879](#)

[Configuring Priority and Preemption for LSPs | 531](#)

refresh-time

Syntax

```
refresh-time seconds;
```

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

Set the refresh time.

Options

seconds—Refresh time.

Range: 1 through 65,535

Default: 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Timers for RSVP Refresh Messages](#) | 851

reliable

Syntax

```
(reliable | no-reliable);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],  
[edit logical-systems logical-system-name protocols rsvp peer-interface peer-interface-name],  
[edit protocols rsvp interface interface-name],  
[edit protocols rsvp peer-interface peer-interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable reliable message delivery on the interface.

To have both refresh reduction and reliable delivery, enable both the **aggregate** and **reliable** statements.

NOTE: For Junos OS Release 16.1 running on MX Series or PTX Series routers, setting **no-reliable** on an interface automatically disables the fast reroute (FRR) scalability enhancements, including refresh reduction, for all label-switched paths (LSPs) traversing the interface.

Default

This option, **reliable**, is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring RSVP Refresh Reduction | 829](#)

[aggregate | 2832](#)

[no-enhanced-frr-bypass | 2876](#)

rsvp

Syntax

```
rsvp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],  
[edit protocols]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Description

Enable Resource Reservation Protocol (RSVP) signaling.

You must include the **rsvp** statement in the configuration to enable RSVP on the router.

The primary purpose of RSVP in Junos OS for EX Series switches is to support dynamic signaling within label switched paths (LSPs).

Default

RSVP is disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Minimum RSVP Configuration | 827](#)

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

rsvp-te (Routing Options)

Syntax

```
rsvp-te entry-name {
  destination-networks network-prefix;
  label-switched-path-template (Multicast) {
    default-template;
    template-name;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-options dynamic-tunnels tunnel-name],
[edit routing-options dynamic-tunnels tunnel-name]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Enable RSVP to automatically establish LSPs for any new PE router added to a full mesh of LSPs. To enable this feature, you must configure the **rsvp-te** statement on all of the PE routers in the full mesh.

Options

entry-name—Specify the entry for the RSVP tunnel.

The other options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring RSVP Automatic Mesh | 850](#)

[Configuring Dynamic Point-to-Multipoint Flooding LSPs](#)

setup-protection

Syntax

```
setup-protection;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Description

The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. You should configure the **setup-protection** statement on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring RSVP Setup Protection](#) | 848

soft-preemption (Protocols RSVP)

Syntax

```
soft-preemption {  
    cleanup-timer seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp preemption],  
[edit protocols rsvp preemption]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable soft preemption to attempt to establish a new path for a preempted LSP before tearing it down.

Options

cleanup-timer—A value of 0 disables soft preemption.

Range: 0 through 10800 seconds

Default: 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS Soft Preemption](#) | 530

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.

Statement introduced in 19.4R1 for cRPD instances.

Description

Configure a static LSP.

Options

lsp-name—Name of the path.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring Static LSPs](#) | 603

subscription

Syntax

```
subscription percentage {
  ct0 percentage;
  ct1 percentage;
  ct2 percentage;
  ct3 percentage;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],
[edit logical-systems logical-system-name protocols rsvp interface interface-name link-protection],
[edit protocols rsvp interface interface-name],
[edit protocols rsvp interface interface-name link-protection]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Configure the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services) or bypass LSP (when you have enabled link protection). **subscription** is the percentage of the link bandwidth that can be used for the RSVP reservation process.

Options

ctnumber percentage—Percentage of the class-type bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs.

Range: 0 through 65,000

Default: 100 percent

percentage—Percentage of the class-type or bypass LSP bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP.

Range: 0 through 65,000

Default: 100 percent

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Bandwidth Subscription Percentage for LSPs | 1210](#)

[Configuring the Amount of Bandwidth Subscribed for Bypass LSPs | 878](#)

traceoptions (Protocols RSVP)

Syntax

```
traceoptions {
  enhanced-frr ;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],
[edit protocols rsvp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable RSVP-level trace options.

Default

The default RSVP-level trace options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

enhanced-frr—(Optional) Enable this option to trace internal events and state changes related to the FRR facility protection enhancements associated with the increased RSVP scaling introduced in Junos OS Release 16.1.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place RSVP tracing output in the file **rsvp-log**.

files number—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **all**—All tracing operations
- **error**—All detected error conditions
- **event**—RSVP-related events
- **io-event [detail] [disable]** —Enable tracing of events that occur within the RSVP I/O task; can only be configured for the master routing instance. The trace output is generally independent of routing instance.
- **io-packets [detail] [disable] [receive] [send]** —Enable tracing of messages as they are received from the network or as they are sent out. This flag can be configured independently for each routing instance. Both bundled and individual messages are identified. Use *detail* to show all objects contained in the message. Use *send* and *receive* to limit tracing to outgoing or incoming packets.
- **lmp**—RSVP-LMP interactions
- **packets**—All RSVP packets
- **path**—All path messages
- **pathtear**—PathTear messages
- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions, including when RSVP-signaled LSPs come up and go down.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Enable only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Tracing RSVP Protocol Traffic](#) | 860

transit

Syntax

```
transit incoming-label {
  bandwidth bps;
  description string;
  link-protection bypass-name name;
  member-interface member-interface;
  next-hop (address | interface-name | address/interface-name);
  node-protection bypass-name name next-next-label label;
  pop;
  stitch {
    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;
  }
  swap out-label;
}
```

Hierarchy Level

```
[edit logical-systems name protocols mpls static-label-switched-path name],
[edit logical-systems name routing-instances name protocols mpls static-label-switched-path name],
[edit logical-systems name tenants name routing-instances name protocols mpls static-label-switched-path name],
[edit protocols mpls static-label-switched-path name],
[edit routing-instances name protocols mpls static-label-switched-path name],
[edit tenants name routing-instances name protocols mpls static-label-switched-path name]
```

Release Information

Statement introduced in Junos OS Release 10.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

switch option introduced in Junos OS Release 14.1X53-D25.

member-interface option introduced in Junos OS Release 17.4R1 for the MX Series and PTX 5000.

Description

Configure a transit static LSP.

NOTE: When configuring transit static LSPs with label operation as **stitch**, the configured next-hop can only be a valid IP address and not an interface name.

Options

incoming-label—Incoming label value.

Range: 1000000 through 1048575

member-interface—Aggregated Ethernet (AE) member interface name.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Static LSPs | 603](#)

[Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP | 688](#)

tunnel-services (RSVP)

Syntax

```
tunnel-services {  
    devices device-names;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp],  
[edit protocols rsvp]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable ultimate-hop popping on point-to-multipoint LSPs. The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically.

Default

Ultimate-hop popping is disabled.

Options

devices *device-names*—Specify which VT interfaces are used to handle the RSVP traffic.

Range: 0 to 8 devices

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs](#) | 859

ultimate-hop-popping

Syntax

```
ultimate-hop-popping;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols mpls],
[edit logical-systems logical-system-name protocols mpls label-switched-path label-switched-path-name],
[edit protocols mpls],
[edit protocols mpls label-switched-path label-switched-path-name]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Description

Enable ultimate-hop popping on LSPs. Configure this statement on the device at the LSP ingress. In ultimate-hop popping, the MPLS label is popped from the IP packet at the PE router. The IP address is checked in a second address lookup (also at the PE router), and then the packet is forwarded to its destination.

Be aware of the following platform requirements and restrictions:

- UHP LSPs using VT interfaces—Supported on all M Series, MX Series, T Series, and TX Matrix routers.
- UHP LSPs using LSI interfaces—Supported on MX 3D Series routers only.
- UHP LSP requirements for the egress PE device—For M Series and T Series routers, a VT interface is needed.
- UHP LSPs and Layer 3 VPNs—UHP LSPs are supported for Layer 3 VPNs configured on MX 3D Series routers only.
- UHP LSPs and VPLS—UHP LSPs are supported for VPLS configured on MX 3D Series routers only. You must configure the *no-tunnel-services* statement at the **[edit routing-instances *routing-instance-name* protocols vpls]** hierarchy level.

Default

Ultimate-hop popping is disabled by default on LSPs. Penultimate-hop popping is the default behavior. In penultimate-hop popping, the final MPLS label is popped from the IP packet at the last provider router in the network before being forwarded to the PE router. The PE router receives the packet and checks the IP address, and then the packet is forwarded to its destination.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Ultimate-Hop Popping for LSPs | 599](#)

[explicit-null | 2611](#)

update-threshold

Syntax

```
update-threshold {  
    threshold-percent;  
    threshold-value threshold-value;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols rsvp interface interface-name],  
[edit protocols rsvp interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

threshold-value option introduced in Junos OS Release 19.4R1 on all platforms.

Description

Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.

Options

threshold-percent—Specify the percentage change in reserved bandwidth to trigger IGP update.

Default: 10 percent

Range: 0.001 through 20 percent

threshold-value *threshold-value*—Specify the change in reserved bandwidth to trigger IGP update. (is capped at 20% of link bandwidth).

If the threshold-value is configured to greater than 20% of bandwidth on that link, the threshold-value is capped at 20% of bandwidth.

For instance, if bandwidth on an interface is 1Gbps, and the **threshold-value** is configured greater than 200Mbps, the **threshold-value** is capped at 200Mbps. The *threshold-percent* is displayed as 20.000% and the **threshold-value** as 200Mbps.

NOTE: The two options, *threshold-percent* and **threshold-value**, are mutually exclusive. You can configure only one option at a given point in time to generate an IGP update for lower bandwidth reservations. As a result, when one option is configured, the other option is calculated and displayed on the CLI.

For instance, on a link of 1Gbps, if the *threshold-percent* is configured to 5%, the **threshold-value** is calculated and displayed as 50Mbps. Similarly, if the **threshold-value** is configured to 50m, then the *threshold-percent* is calculated and displayed as 5%.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the RSVP Update Threshold on an Interface](#) | 833

LDP Configuration Statements

IN THIS CHAPTER

- [allow-subnet-mismatch](#) | 2908
- [authentication-algorithm](#) | 2909
- [authentication-key \(Protocols LDP\)](#) | 2912
- [authentication-key-chain \(Protocols LDP\)](#) | 2913
- [auto-targeted-session](#) | 2915
- [bfd-liveness-detection \(Protocols LDP\)](#) | 2917
- [deaggregate](#) | 2919
- [disable \(Protocols LDP\)](#) | 2920
- [dod-request-policy](#) | 2921
- [downstream-on-demand](#) | 2922
- [ecmp](#) | 2923
- [egress-policy](#) | 2924
- [explicit-null \(Protocols LDP\)](#) | 2925
- [export \(Protocols LDP\)](#) | 2926
- [failure-action \(Protocols LDP\)](#) | 2927
- [fec](#) | 2928
- [graceful-restart \(Protocols LDP\)](#) | 2930
- [hello-interval \(Protocols LDP\)](#) | 2931
- [helper-disable \(LDP\)](#) | 2932
- [holddown-interval](#) | 2933
- [hold-time \(Protocols LDP\)](#) | 2934
- [ignore-lsp-metrics](#) | 2935
- [igp-synchronization](#) | 2936
- [import \(Protocols LDP\)](#) | 2937
- [ingress-policy](#) | 2938
- [interface \(Protocols LDP\)](#) | 2939
- [keepalive-interval](#) | 2940
- [keepalive-timeout](#) | 2941

- l2-smart-policy | 2942
- label-withdrawal-delay | 2943
- ldp | 2944
- ldp-synchronization | 2946
- log-updown (Protocols LDP) | 2947
- make-before-break (LDP) | 2948
- mapping-server-entry | 2949
- maximum-neighbor-recovery-time | 2950
- mldp-inband-signalling (Protocols Multipoint LDP) | 2951
- mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain) | 2952
- mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain) | 2954
- mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain) | 2955
- mofrr-primary-path-selection-by-routing (Multicast-Only Fast Reroute) | 2956
- neighbor (Protocols LDP) | 2958
- no-forwarding | 2959
- oam (Protocols LDP) | 2960
- p2mp (Protocols LDP) | 2962
- p2mp-ldp-next-hop | 2964
- periodic-traceroute | 2965
- policing (Protocols LDP) | 2967
- policy (Multicast-Only Fast Reroute) | 2968
- policy (Protocols Multipoint LDP) | 2971
- preference (Protocols LDP) | 2972
- prefix-segment (Routing Options) | 2973
- prefix-segment-range | 2974
- reconnect-time | 2976
- recovery-time | 2977
- session (Protocols LDP) | 2978
- session-group | 2980
- session-protection | 2982
- source-packet-routing | 2983
- stream-protection (Multicast-Only Fast Reroute) | 2984
- strict-targeted-hellos | 2985
- targeted-hello | 2986

- [traceoptions \(Protocols LDP\) | 2987](#)
- [track-igp-metric | 2990](#)
- [track-igp-metric \(LSP\) | 2991](#)
- [traffic-statistics \(Protocols LDP\) | 2992](#)
- [transport-address | 2994](#)
- [version \(BFD\) | 2996](#)

allow-subnet-mismatch

Syntax

```
allow-subnet-mismatch;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp interface interface-name],  
[edit protocols ldp interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.

Default

The source address in the LDP link hello packet is matched against the interface address.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Ignoring the LDP Subnet Check | 1122](#)

authentication-algorithm

Syntax

```
authentication-algorithm algorithm;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name protocols ldp session session-address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name
  neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp session session-address],
[edit logical-systems logical-system-name routing-options bmp],
[edit logical-systems logical-system-name routing-options bmp station station-name],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit protocols ldp session session-address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols ldp session session-address],
[edit routing-options bmp],
[edit routing-options bmp station station-name]
```

Release Information

Statement introduced in Junos OS Release 7.6.

Statement introduced for BGP in Junos OS Release 8.0.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced for BMP in Junos OS Release 13.3.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure an authentication algorithm type.

NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha-256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.
- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as hmac-sha-256-128 and hmac-md5-96 on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as hmac-md5-96 on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.
- When you configure two IPsec proposals at both ends of a tunnel, such as the **authentication-algorithm hmac-sha-256-128** and **authentication-algorithm hmac-md5-96** statements at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the **authentication-algorithm hmac-md5-96** and **authentication-algorithm hmac-sha-256-128** statements at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is hmac-md5-96 and not the stronger algorithm of hmac-sha-256-128. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the 3des-cbc algorithm is chosen and not the aes-cfb algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, hmac-sha-256-128 is selected as the authentication method.
- You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.

Options

algorithm—Specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

Default: **hmac-sha-1-96**

NOTE: The default is not displayed in the output of the **show bgp bmp** command unless a key or key-chain is also configured.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Router 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](#) | 1118

authentication-key-chain (Protocols LDP)

Syntax

```
authentication-key-chain key-chain;
```

Hierarchy Level

```
[edit logical-systems name protocols ldp session address],  
[edit logical-systems name routing-instances instance-name protocols ldp session address],  
[edit protocols ldp session address],  
[edit routing-instances instance-name protocols ldp session address]
```

Release Information

Statement introduced in Junos OS Release 8.0.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the **0.0.0.0/allow** statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.

NOTE: You must also configure an authentication algorithm using the **authentication-algorithm** statement.

Options

key-chain—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

Configuring Miscellaneous LDP Properties | 1115

authentication-algorithm | 2909

auto-targeted-session

Syntax

```
auto-targeted-session {
  maximum-sessions seconds;
  teardown-delay seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Configure session parameters for LDP sessions established with the remote LFA node that are automatically targeted using the loopback addresses. Configure parameters of automatically targeted sessions for remote LFA only.

Options

maximum-sessions *seconds* —Specify the maximum number of auto-targeted LDP sessions allowed. Include this statement to optimize the use of router memory.

Default: 100

Range: 1 through 1000

teardown-delay *seconds* —Specify the minimum time period for which the auto-targeted session must be alive before tearing down the auto-targeted LDP sessions to the remote LFA node. Include this statement to prevent rapid route-resolution in case of temporary change in IGP topology.

Default: 90 seconds

Range: 1 through 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

no-eligible-remote-backup

| *remote-backup-calculation*

bfd-liveness-detection (Protocols LDP)

Syntax

```

bfd-liveness-detection {
  detection-time threshold milliseconds;
  ecmp;
  failure-action {
    remove-nexthop;
    remove-route;
  }
  holddown-interval seconds;
  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);
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols ldp oam],
[edit logical-systems logical-system-name protocols ldp oam fec address],
[edit protocols ldp oam],
[edit protocols ldp oam fec address]

```

Release Information

Statement introduced in Junos OS Release 7.6.

Support for the **bfd-liveness-detection** statement at the **[edit protocols ldp oam fec address]** hierarchy level and the **ecmp** option added in Junos OS Release 9.0.

Support for the **failure-action** statement with the **remove-nexthop** and **remove-route** options and the **holddown-interval** statement added in Junos OS Release 9.4.

Description

Enable Bidirectional Forwarding Detection (BFD) for all MPLS LSPs or for just a specific LSP.

Options

minimum-interval—Minimum transmit and receive interval.

Range: 50 through 255,000 milliseconds

Default: 50

minimum-receive-interval—Minimum receive interval.

Range: 50 through 255,000 milliseconds

Default: 50

minimum-transmit-interval—Minimum transmit interval.

Range: 50 through 255,000 milliseconds

Default: 50

multiplier—Detection time multiplier.

Range: 50 through 255

Default: 3

The other options are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring BFD for LDP LSPs](#) | 975

deaggregate

Syntax

```
deaggregate | no-deaggregate;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Control forwarding equivalence class (FEC) deaggregation on the router. The use of the **deaggregate** statement in LDP is a standard practice that we recommend for LDP deployments.

Default

Deaggregation is disabled on the router.

Options

deaggregate—Deaggregate FECs.

no-deaggregate—Aggregate FECs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring FEC Deaggregation](#) | 973

disable (Protocols LDP)

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp interface
  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options graceful-restart],
[edit protocols ldp graceful-restart],
[edit protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp interface interface-name],
[edit routing-instances routing-instance-name routing-options graceful-restart]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.

Default

LDP is enabled on interfaces configured with the LDP **interface** statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the **[edit routing-options]** hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Enabling and Disabling LDP | 934](#)

[Configuring LDP Graceful Restart | 960](#)

dod-request-policy

Syntax

```
dod-request-policy dod-request-policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the name of the LDP downstream on demand request policy. The **dod-request-policy** statement performs exact match, as a result, LDP sends label request messages only for those FECs matching in the downstream on demand request policy.

Options

dod-request-policy-name—Specify the name of the downstream on demand request policy.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring LDP Downstream on Demand](#) | 1046

downstream-on-demand

Syntax

```
downstream-on-demand;
```

Hierarchy Level

```
[edit logical systems logical-system-name protocols ldp session session-address],  
[edit protocols ldp session session-address]
```

Release Information

Statement introduced in Junos OS Release 12.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring LDP Downstream on Demand](#) | 1046

ecmp

Syntax

```
ecmp;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-liveness-detection],  
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-liveness-detection],  
[edit protocols ldp oam bfd-liveness-detection],  
[edit protocols ldp oam fec address bfd-liveness-detection]
```

Release Information

Statement introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.

Description

Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the **ecmp** statement, you must also configure the **periodic-traceroute** statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the **periodic-traceroute** statement at the global hierarchy level ([**edit protocols ldp oam**]) while only configuring the **ecmp** statement for a specific FEC ([**edit protocols ldp oam fec *address* bfd-liveness-detection**]).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring ECMP-Aware BFD for LDP LSPs](#) | 979

egress-policy

Syntax

```
egress-policy [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Control the prefixes advertised into LDP.

Default

Only the loopback address is advertised.

Options

policy-names—Name of one or more routing policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Prefixes Advertised into LDP from the Routing Table](#) | 971

explicit-null (Protocols LDP)

Syntax

```
explicit-null;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Advertise label 0 to the egress router of a label-switched path (LSP).

Default

If you do not include the **explicit-null** statement in the MPLS configuration, label 3 (implicit null) is advertised.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router](#) | 1117

export (Protocols LDP)

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options

policy-names—Name of one or more routing policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Filtering Outbound LDP Label Bindings](#) | 965

failure-action (Protocols LDP)

Syntax

```
failure-action {  
    remove-nexthop;  
    remove-route;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-livenesss-detection],  
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-livenesss-detection],  
[edit protocols ldp oam bfd-livenesss-detection],  
[edit protocols ldp oam fec address bfd-livenesss-detection]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.

Options

remove-nexthop—Remove a route corresponding to a next hop of the LSP's route at the ingress node when a BFD session failure event is detected.

remove-route—Remove the route corresponding to an LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring a Failure Action for the BFD Session on an LDP LSP](#) | 980

fec

Syntax

```
fec fec-address {
  bfd-liveness-detection {
    detection-time threshold milliseconds;
    ecmp;
    failure-action {
      remove-nexthop;
      remove-route;
    }
    holddown-interval milliseconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (0 | 1 | automatic);
  }
  no-bfd-liveness-detection;
  periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-systems-name protocols ldp oam],
[edit protocols ldp oam]
```


Release Information

Statement introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 12.2 for EX Series switches.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).

Options

fec-address—Specify the FEC address.

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring BFD for LDP LSPs](#) | 975

graceful-restart (Protocols LDP)

Syntax

```
graceful-restart {  
    disable;  
    helper-disable;  
    maximum-neighbor-recovery-time value;  
    reconnect-time seconds;  
    recovery-time value;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

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](#) | 960

hello-interval (Protocols LDP)

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name protocols ldp targeted-hello],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp interface
  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp targeted-hello],
[edit protocols ldp interface interface-name],
[edit protocols ldp targeted-hello],
[edit routing-instances routing-instance-name protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp targeted-hello]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for LDP targeted hellos added in Junos OS Release 9.5.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the **hello-interval** statement.

Options

seconds—Length of time between transmission of hello packets.

Range: 1 through 65,535 seconds

Default: 5 seconds for link hello messages, 15 seconds for targeted hello messages

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the LDP Timer for Hello Messages](#) | 934

helper-disable (LDP)

Syntax

```
helper-disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp graceful-restart],  
[edit protocols ldp graceful-restart],  
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.

Default

Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring LDP Graceful Restart](#) | 960

holddown-interval

Syntax

```
holddown-interval holddown-interval;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-livenesss-detection],  
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-livenesss-detection],  
[edit protocols ldp oam bfd-livenesss-detection],  
[edit protocols ldp oam fec address bfd-livenesss-detection]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Description

Specify how long the BFD session should be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

Options

holddown-interval—Number of seconds the BFD session should remain up before adding the route or next hop.

Default: 0 seconds

Range: 0 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Holddown Interval for the BFD Session](#) | 980

hold-time (Protocols LDP)

Syntax

```
hold-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name protocols ldp targeted-hello],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp interface
  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp targeted-hello],
[edit protocols ldp interface interface-name],
[edit protocols ldp targeted-hello],
[edit routing-instances routing-instance-name protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp targeted-hello]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for LDP targeted hellos added in Junos OS Release 9.5.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the **hold-time** statement.

Options

seconds—Hold-time value.

Range: 1 through 65,535 seconds

Default: 15 seconds for link hello messages, 45 seconds for targeted hello messages

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Delay Before LDP Neighbors Are Considered Down](#) | 935

ignore-lsp-metrics

Syntax

```
ignore-lsp-metrics;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf traffic-engineering shortcuts],  
[edit protocols ospf traffic-engineering shortcuts]
```

Release Information

Statement introduced in Junos OS Release 7.5.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Cause OSPF to ignore the RSVP LSP metric.

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks](#) | 1118

igp-synchronization

Syntax

```
igp-synchronization holddown-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

Options

holddown-interval *seconds*—Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational.

Default: 10 seconds

Range: 10 through 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LDP Synchronization with the IGP on the Router](#) | 1121

import (Protocols LDP)

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options

policy-names—Name of one or more routing policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Filtering Inbound LDP Label Bindings](#) | 963

ingress-policy

Syntax

```
ingress-policy [ ingress-policy-names ];
```

Hierarchy Level

```
[edit logical-system logical-system-name protocols ldp entropy-label],  
[edit logical-system logical-system-name protocols ldp oam],  
[edit protocols ldp entropy-label],  
[edit protocols ldp oam]
```

Release Information

Statement introduced in Junos OS Release 9.4.

Statement introduced at the **[edit protocols ldp entropy-label]** hierarchy level in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Configure an LDP ingress policy for either the entropy label or Operation, Administration, and Management (OAM).

For OAM, configure the ingress policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under **[edit protocols ldp oam bfd-liveness-detection]** are applied.

Options

ingress-policy-names—Specify the names of the ingress policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring OAM Ingress Policies for LDP | 1238](#)

[Configuring the Entropy Label for LSPs | 573](#)

interface (Protocols LDP)

Syntax

```
interface interface-name {  
    disable;  
    hello-interval seconds;  
    hold-time seconds;  
    transport-address (interface | loopback);  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable LDP on one or more router interfaces.

Default

LDP is disabled on all interfaces.

Options

interface-name—Name of an interface. To configure all interfaces, specify **all**.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Enabling and Disabling LDP](#) | 934

keepalive-interval

Syntax

```
keepalive-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Set the keepalive interval value.

Options

seconds—Keepalive value.

Range: 1 through 65,535

Default: 10 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Interval for LDP Keepalive Messages](#) | 938

keepalive-timeout

Syntax

```
keepalive-timeout seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.

Options

seconds—Keepalive timeout value.

Range: 1 through 65,535

Default: 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the LDP Keepalive Timeout](#) | 938

I2-smart-policy

Syntax

```
I2-smart-policy;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring LDP IPv4 FEC Filtering](#) | 975

label-withdrawal-delay

Syntax

```
label-withdrawal-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Delay the withdrawal of labels to reduce router workload during IGP convergence.

Options

seconds—Configure the number of seconds to wait before withdrawing labels for the LDP LSPs.

Default: 60 seconds

Range: 0 through 300 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Label Withdrawal Timer](#) | 1121

Ldp

Syntax

```

ldp {
  auto-targeted-session;
  (deaggregate | no-deaggregate);
  dual-transport ;
  egress-policy [ policy-names ];
  entropy-label;
  explicit-null;
  export [ policy-names ];
  family (Protocols LDP);
  graceful-restart;
  igp-synchronization;
  import [ policy-names];
  interface (interface-name | all);
  keepalive-interval seconds;
  keepalive-timeout seconds;
  log-updown;
  longest-match;
  make-before-break;
  neighbor;
  no-forwarding;
  oam;
  p2mp;
  policing;
  preference preference;
  session;
  session-group;
  session-protection;
  strict-targeted-hellos;
  traceoptions;
  track-igp-metric;
  traffic-statistics;
  transport-address (address | interface | router-id);
  transport-preference [ipv4 | ipv6];
}

```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
 [edit protocols],

[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

dual-transport statement introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

family statement introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

transport-preference option introduced in Junos OS Release 16.1 for the M320 Series, MX Series, and PTX Series.

Description

Enable LDP routing on the router or switch.

You must include the **ldp** statement in the configuration to enable LDP on the router or switch.

Default

LDP is disabled on the router.

Options

transport-preference ipv4 | ipv6— Select the preferred transport for TCP connection when both IPv4 and IPv6 are enabled. If **transport-preference ipv4** is configured, LDP will attempt to establish the TCP connection using IPv4. If **transport-preference ipv6** is configured, LDP will attempt to establish the TCP connection using IPv6.

Default: ipv6

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Minimum LDP Configuration | 933](#)

[Enabling and Disabling LDP | 934](#)

Ldp-synchronization

Syntax

```
ldp-synchronization {  
    disable;  
    hold-time seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf interface interface-name],  
[edit logical-systems logical-system-name routing-instances routing-instance-name  
protocols ospf interface interface-name],  
[edit protocols ospf interface interface-name],  
[edit routing-instances routing-instance-name protocols ospf interface interface-name]
```

Release Information

Statement introduced in Junos OS Release 7.5.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.

Options

The other statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LDP Synchronization with the IGP on LDP Links](#) | 1121

log-updown (Protocols LDP)

Syntax

```
log-updown {  
    trap disable;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Disable LDP traps on the router, logical system, or routing instance.

Options

trap disable—Disable LDP traps.

Default: LDP traps are enabled on the router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling SNMP Traps for LDP](#) | 1120

make-before-break (LDP)

Syntax

```
make-before-break {
    timeout seconds;
    switchover-delay seconds;
}
```

Hierarchy Level

```
[edit protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Description

Configures make before break (MBB) for multicast LDP (MLDP) link protection to ensure minimum packet loss when attempting to signal a new label-switched path (LSP) before tearing down the old LSP path.

Options

timeout *seconds*—Specify a value to change a make -before-break timeout for point-to-multipoint LSPs.

Even if an MBB acknowledgment is not received for a point-to-multipoint LSP before the specified timeout period expires, the label-switching router (LSR) performs an MBB switchover from the old LSR to the new upstream LSR.

Range: 1 through 300 seconds

Default: 30 seconds

switchover-delay *seconds*—Specify a value to change switchover delay for a point-to-multipoint LSP from the old LSR to the new upstream LSR. If an MBB acknowledgment is received on a point of local repair (PLR) router, the PLR waits for the specified seconds to switch its upstream LSR from the old LSR to the new LSR.

Range: 1 through 300 seconds

Default: 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

mapping-server-entry

Syntax

```
mapping-server-entry mapping-server-name {  
    prefix-segment prefix;  
    prefix-segment-range prefix-segment-range-name;  
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing],  
[edit logical-systems name routing-options source-packet-routing],  
[edit routing-instances name routing-options source-packet-routing],  
[edit routing-options source-packet-routing]
```

Release Information

Statement introduced in Junos OS Release 18.2R1.

Description

Configure an LDP mapping server to enable interoperability between islands of devices supporting only segment routing and only LDP in an LDP network domain.

The mapping server configuration can be included on any device in the segment routing network.

Options

mapping-server-entry-name—Name of the LDP mapping server.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing

maximum-neighbor-recovery-time

Syntax

```
maximum-neighbor-recovery-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp graceful-restart],
[edit protocols ldp graceful-restart],
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Release Information

Statement introduced before Junos OS Release 7.4. Statement changed from **maximum-recovery-time** to **maximum-neighbor-recovery-time** in Junos OS Release 9.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.

Options

seconds—Configure the maximum recovery time, in seconds.

Range: 120 through 1800 seconds

Default: 140 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Recovery Time and Maximum Recovery Time | 962](#)

[Configuring Graceful Restart Options for LDP](#)

[no-strict-lsa-checking](#)

[recovery-time](#)

mldp-inband-signalling (Protocols Multipoint LDP)

Syntax

```
mldp-inband-signalling {
  policy policy-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim],
[edit protocols pim],
```

Release Information

Statement introduced in Junos OS Release 13.2.

Support added in Junos OS Release 18.2R1 for using this command in conjunction with distributed MLD or distributed IGMP.

Description

Multipoint LDP (mLDP) in-band signalling lets you carry multicast traffic across an existing IP/MPLS backbone, while avoiding the use of PIM in the provider core.

On the label-edge router (LER), enable PIM to use mLDP in-band signaling for the upstream neighbors when the LER does not detect a PIM upstream neighbor. On the egress nodes, configure the MPLS LSP root in the PIM configuration, using the **policy** statement.

When used in conjunction with distributed MLD or distributed IGMP, mLDP inband signalling supports interconnecting separate PIM domains via a MPLS-based core. To enable the inter-working, **chassis network-services enhanced-ip** must be enabled and you need to set the **dynamic-profiles profile-name protocols igmp|mld interface interface-name** to **distributed**. Enabling this command, **mldp-inband-signalling**, has PIM act as a multipoint LDP inband edge router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs](#) | 1072

mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)

Syntax

```
mofrr-asm-starg;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast
  stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.4R1 for QFX Series switches.

Description

Enable **mofrr-asm-starg** to include any-source multicast (ASM) for (*,G) joins in the Multicast-only fast reroute (MoFRR).

NOTE: **mofrr-asm-starg** applies to IP-PIM only. When enabled for group G, *,G will undergo MoFRR as long as there is no S#,G for Group G. In other words, *,G MoFRR will cease and any old states will be torn down when S#,G is created. Note too, that **mofrr-asm-starg** is not supported for mLDP (since mLDP itself does not support *,G).

In a PIM domain with MoFRR enabled, the default for **stream-protection** is S,G routes only.

Context: Multicast-only fast reroute (MoFRR) can be used to reduce traffic loss in a multicast distribution tree in the event of link down. To employ MoFRR, a downstream router is configured with an alternative path back towards the source, over which it receives a backup live stream of the same multicast traffic. That router propagates the same (S,G) join toward both upstream neighbors in order to create duplicate multicast trees. If a failure is detected on the primary tree, the router switches to the backup tree to prevent packet loss.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

Understanding Multicast-Only Fast Reroute on Switches

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)

Syntax

```
mofrr-disjoint-upstream-only;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast
stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.4R1 for QFX Series switches.

Description

When you configure multicast-only fast reroute (MoFRR) in a PIM domain, allow only a disjoint RPF (an RPF on a separate plane) to be selected as the backup RPF path.

In a multipoint LDP MoFRR domain, the same label is shared between parallel links to the same upstream neighbor. This is not the case in a PIM domain, where each link forms a neighbor. The

mofrr-disjoint-upstream-only statement does not allow a backup RPF path to be selected if the path goes to the same upstream neighbor as that of the primary RPF path. This ensures that MoFRR is triggered only on a topology that has multiple RPF upstream neighbors.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

[Understanding Multicast-Only Fast Reroute on Switches](#)

[Example: Configuring Multicast-Only Fast Reroute in a PIM Domain](#)

[Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches](#)

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)

Syntax

```
mofrr-no-backup-join;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast
stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.4R1 for QFX Series switches.

Description

When you configure multicast-only fast reroute (MoFRR) in a PIM domain, prevent sending join messages on the backup path, but retain all other MoFRR functionality.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

Understanding Multicast-Only Fast Reroute on Switches

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

mofr-primary-path-selection-by-routing (Multicast-Only Fast Reroute)

Syntax

```
mofr-primary-path-selection-by-routing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast
stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.4R1 for QFX Series switches.

Description

MoFRR is supported on both equal-cost multipath (ECMP) paths and non-ECMP paths. Unicast loop-free alternate (LFA) routes need to be enabled to support MoFRR on non-ECMP paths. LFA routes are enabled with the link-protection statement in the interior gateway protocol (IGP) configuration. When you enable link protection on an OSPF or IS-IS interface, Junos OS creates a backup LFA path to the primary next hop for all destination routes that traverse the protected interface.

In the context of load balancing, MoFRR prioritizes the disjoint backup in favor of load balancing the available paths.

For Junos OS releases before 15.1R7, for both ECMP and Non-ECMP scenarios, the default MoFRR behavior was *sticky*, that is, if the Active link went down, the Active Path selection would give preference to Backup Path for the transition. The Active Path would not follow the unicast selected gateway

Starting in Junos OS Release 15.1R7 however, the default behavior for non-EMCP scenarios is to be *nonsticky*, that is, the selection of Active Path strictly follows unicast selected gateway. MoFRR no longer chooses a unicast LFA path to become the MoFRR Active path; only a unicast LFA path can be selected to become MoFRR Backup.

Default

By default, the backup path gets promoted to be the primary path when MoFRR is configured in a PIM domain.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

Understanding Multicast-Only Fast Reroute on Switches

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

neighbor (Protocols LDP)

Syntax

```
neighbor neighbor-address;
```

Hierarchy Level

```
[edit protocols ldp],
```

Release Information

Statement introduced in Junos OS Release 14.2R1 on all platforms.

Description

Configure an LDP targeted neighbor.

LDP sends a targeted hello message to the configured remote neighbor with a targeted request-send-targeted hello message (T) bit set. If the remote neighbor allows receipt of asymmetric hello messages, or if it is configured with the source address as the targeted neighbor, it responds with a targeted hello message. The receipt of a targeted hello message establishes a targeted adjacency with the remote neighbor as described in RFC 5036. Subsequently, a targeted session is established to the remote neighbor.

Options

neighbor-address—IP address of the remote LDP neighbor.

Required Privilege Level

routing

RELATED DOCUMENTATION

[Minimum LDP Configuration | 933](#)

[Enabling and Disabling LDP | 934](#)

no-forwarding

Syntax

```
no-forwarding;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Do not add ingress routes to the inet.0 routing table even if [traffic-engineering](#) **bgp-igp** (configured at the [\[edit protocols mpls\]](#) hierarchy level) is enabled.

Default

The **no-forwarding** statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when [traffic-engineering](#) **bgp-igp** is enabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Preventing Addition of Ingress Routes to the inet.0 Routing Table](#) | 1116

[Configuring Virtual-Router Routing Instances in VPNs](#)

oam (Protocols LDP)

Syntax

```

oam {
  bfd-liveness-detection {
    detection-time threshold milliseconds;
    ecmp;
    failure-action {
      remove-nexthop;
      remove-route;
    }
    holddown-interval milliseconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (0 | 1 | automatic);
  }
  fec fec-address;
  ingress-policy ingress-policy-name;
  lsp-ping-interval seconds;
  periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
  }
}

```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp]


```
[edit protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 7.6.

lsp-ping-interval option introduced in Junos OS Release 9.4.

Description

Configure Operation, Administration, and Maintenance (OAM) and Bidirectional Forwarding Detection (BFD) protocol for LDP.

Options

fec fec-address—Specify the forwarding equivalence class (FEC) address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.

lsp-ping-interval seconds—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command.

Default: 60 seconds

Range: 30 through 3,600 seconds

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring BFD for LDP LSPs](#) | 975

p2mp (Protocols LDP)

Syntax

```
p2mp {
  no-rsvp-tunneling;
  recursive;
  root-address root-address;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 11.2.

no-rsvp-tunneling option added in Junos OS Release 16.1R5.

Description

Enable point-to-multipoint MPLS LSPs in an LDP-signaled LSP.

Options

no-rsvp-tunneling—(Optional) Disable LDP point-to-multipoint LSPs from using RSVP-TE LSPs for tunneling, and use LDP paths instead.

NOTE: The **no-rsvp-tunneling** option is introduced in Junos OS Release 16.1R5, 17.3R1, 17.2R2, 16.2R3, and later releases.

Starting in Junos OS Release 12.3R1, Junos OS provides support for Multipoint LDP (M-LDP) for Targeted LDP (T-LDP) sessions with unicast replication, in addition to link sessions. As a result, the default behavior of M-LDP over RSVP tunneling is similar to unicast LDP. However, because T-LDP is chosen over LDP and link sessions to signal point-to-multipoint LSPs, the **no-rsvp-tunneling** option enables LDP natively throughout the network.

recursive—(Optional) Configure point-to-multipoint recursive parameters, including **route**.

root-address *root-address*—(Optional) Specify the root address of the point-to-multipoint LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs

[Point-to-Multipoint LSPs Overview](#) | 691

p2mp-ldp-next-hop

Syntax

```
p2mp-ldp-next-hop {
  root-address root-address{
    lsp-id id;
  }
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options static route
destination-prefix],
[edit logical-systems logical-system-name routing-options static route destination-prefix],
[edit routing-instances routing-instance-name routing-options static route destination-prefix].
[edit routing-options static route destination-prefix]
```

Release Information

Statement introduced in Junos OS Release 13.3.

Description

Specify a point-to-multipoint LDP label-switched path (LSP) as the next hop for a static route, and configure a root and provide an *lsp-id* on that LDP-signalled label-switched path.

Options

root-address *root address*— Specify the root address of the point-to-multipoint LSP.

lsp-id *id*— Specify the generic LSP identifier. The range is 1 through 65535.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

periodic-traceroute

Syntax

```
periodic-traceroute {
  disable;
  exp exp-value;
  fanout fanout-value;
  frequency minutes;
  paths number-of-paths;
  retries retry-attempts;
  source address;
  ttl ttl-value;
  wait seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam],
[edit logical-systems logical-system-name protocols ldp oam fec fec-address],
[edit protocols ldp oam],
[edit protocols ldp oam fec fec-address]
```

Release Information

Statement introduced in Junos OS Release 8.4.

Support added at the **[edit protocols ldp oam]** and **[edit logical-systems *logical-system-name* protocols ldp oam]** hierarchy levels in Junos OS Release 9.0.

Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description

Enable tracing of forwarding equivalence classes (FECs) for LDP LSPs.

Options

disable—(Optional) Disable tracing for a specific FEC. This option is available at the **[edit protocols ldp oam fec *fec-address* periodic-traceroute]** and **[edit logical-systems *logical-system-name* protocols ldp oam fec *fec-address* periodic-traceroute]** hierarchy levels only.

exp *exp-value*—(Optional) Specify the class of service to use when sending probes.

Default: 7

Range: 0 through 7

fanout *fanout-value*—(Optional) Specify the maximum number of next hops to search per node.

Default: 16

Range: 1 through 16

frequency *minutes*—(Optional) Specify the interval between traceroute attempts.

Default: 60 minutes

Range: 15 through 120 minutes

paths *number-of-paths*—(Optional) Specify the maximum number of paths to search.

Default: 3

Range: 1 through 255

retries *retry-attempts*—(Optional) Specify the number of attempts to send a probe to a specific node before giving up.

Default: 3

Range: 1 through 9

source *address*—(Optional) Specify the IPv4 source address to use when sending probes.

ttl *value*—(Optional) Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.

Default: 64

Range: 1 through 255

wait *seconds*—(Optional) Specify the wait interval before resending a probe packet.

Default: 10 seconds

Range: 5 through 15 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LDP LSP Traceroute](#) | **1123**

policing (Protocols LDP)

Syntax

```
policing {  
  fec fec-address {  
    ingress-traffic filter-name;  
    transit-traffic filter-name;  
  }  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Enable policing of forwarding equivalence classes (FECs) for LDP.

Options

fec *fec-address*—Specify the address for the FEC.

ingress-traffic *filter-name*—Specify the name of the filter for policing ingress FEC traffic.

transit-traffic *filter-name*—Specify the name of the filter for policing transit FEC traffic.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Policers for LDP FECs | 974](#)

policy (Multicast-Only Fast Reroute)

Syntax

```
policy policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast
stream-protection],
[edit logical-systems logical-system-name routing-options multicast stream-protection],
[edit routing-instances routing-instance-name routing-options multicast stream-protection],
[edit routing-options multicast stream-protection]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.4R1 for QFX Series switches.

Description

When you configure multicast-only fast reroute (MoFRR), apply a routing policy that filters for a restricted set of multicast streams to be affected by your MoFRR configuration. You can apply filters that are based on source or group addresses.

For example:

```
routing-options {
  multicast {
    stream-protection {
      policy mofrr-select;
    }
  }
}
policy-statement mofrr-select {
  term A {
    from {
      source-address-filter 225.1.1.1/32 exact;
    }
    then {
      accept;
    }
  }
  term B {
    from {
```



```

        source-address-filter 226.0.0.0/8 orlonger;
    }
    then {
        accept;
    }
}
term C {
    from {
        source-address-filter 227.1.1.0/24 orlonger;
        source-address-filter 227.4.1.0/24 orlonger;
        source-address-filter 227.16.1.0/24 orlonger;
    }
    then {
        accept;
    }
}
term D {
    from {
        source-address-filter 227.1.1.1/32 exact;
    }
    then {
        reject; #MoFRR disabled
    }
}
term E {
    from {
        route-filter 227.1.1.0/24 orlonger;
    }
    then accept;
}
...
}

```

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

Understanding Multicast-Only Fast Reroute on Switches

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

policy (Protocols Multipoint LDP)

Syntax

```
policy policy-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols pim mldp-inband-signalling],  
[edit protocols pim mldp-inband-signalling]
```

Release Information

Statement introduced in Junos OS Release 13.2.

Description

Multipoint LDP (M-LDP) in-band signaling enables you to carry multicast traffic across an existing IP/MPLS backbone, while avoiding the use of PIM in the provider core.

On the egress nodes of the point-to-multipoint LSP, specify an M-LDP join translation filter policy where PIM messages are translated into M-LDP FEC bindings. The policy statement is needed when internal BGP (IBGP) is not available in the core site or to override IBGP-based LSP root detection.

The filter policy is configured at the **[edit policy-options]** hierarchy level. The policy generally specifies one or more source-address filters and the point-to-multipoint LDP root IP address using the **p2mp-lsp-root** policy action.

Options

policy-name—Name of a policy configured at the **[edit policy-options]** hierarchy level.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs](#) | 1072

preference (Protocols LDP)

Syntax

```
preference preference;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit protocols ldp interface interface-name],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit protocols ldp interface interface-name],  
[edit routing-instances routing-instance-name protocols ldp interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Set the route preference level for LDP routes.

Options

preference—Preferred value.

Range: 0 through 255

Default: 9

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LDP Route Preferences](#) | 959

prefix-segment (Routing Options)

Syntax

```
prefix-segment prefix-segment {
    index index;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing mapping-server-entry],
[edit logical-systems name routing-options source-packet-routing mapping-server-entry],
[edit routing-instances name routing-options source-packet-routing mapping-server-entry],
[edit routing-options source-packet-routing mapping-server-entry]
```

Release Information

Statement introduced in Junos OS Release 18.2R1.

Description

Configure the IP address and index number of the prefix segment for the LDP mapping server.

Options

prefix-segment—IP address of the prefix segment.

index index—Prefix segment index.

Range: 0 through 199999

Required Privilege Level

routing

RELATED DOCUMENTATION

[LDP Mapping Server for Interoperability of Segment Routing with LDP Overview | 1110](#)
[source-packet-routing | 2983](#)

prefix-segment-range

Syntax

```
prefix-segment-range prefix-segment-range-name {
    attached;
    domain-wide-flooding;
    no-node-segment;
    size size;
    start-index start-index;
    start-prefix start-prefix;
}
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options source-packet-routing mapping-server-entry],
[edit logical-systems name routing-options source-packet-routing mapping-server-entry],
[edit routing-instances name routing-options source-packet-routing mapping-server-entry],
[edit routing-options source-packet-routing mapping-server-entry]
```

Release Information

Statement introduced in Junos OS Release 18.2R1.

attached, **domain-wide-flooding**, and **no-node-segment** options introduced in Junos OS Release 19.1R1.

Description

Configure the prefix segment range for the LDP mapping server.

Options

prefix-segment-range-name—Name of the prefix segment range.

attached—(Optional) Set the flag in IS-IS mapping server advertisement to indicate that the prefixes and SIDs advertised in the SID or Label Binding TLV are directly connected to their originators.

domain-wide-flooding—(Optional) Set an S flag in the IS-IS mapping server advertisement to indicate that the SID or Label Binding TLV is flooded across the entire routing domain.

no-node-segment—(Optional) Clear the node segment flag in the mapping server prefix segment to indicate that the prefix has originated from a single node.

size *size*—Size of prefix segment range.

Range: 1 through 1024

start-index *start-index*—Include start index.

Range: 0 through 199999

start-prefix *start-prefix*—Include start prefix.

Required Privilege Level

- routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[LDP Mapping Server for Interoperability of Segment Routing with LDP Overview | 1110](#)

[source-packet-routing | 2983](#)

[mapping-server-entry | 2949](#)

reconnect-time

Syntax

```
reconnect-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],  
[edit protocols ldp graceful-restart],  
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Release Information

Statement introduced in Junos OS Release 9.1.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.

Options

seconds—Time required for reconnection.

Range: 30 through 300

Default: 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring LDP Graceful Restart](#) | [960](#) on *MPLS Applications User Guide*

Configuring Graceful Restart Options for LDP

recovery-time

Syntax

```
recovery-time seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp graceful-restart],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp graceful-restart],  
[edit protocols ldp graceful-restart],  
[edit routing-instances routing-instance-name protocols ldp graceful-restart]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the amount of time a router waits for LDP to restart gracefully.

Options

seconds—Configure the recovery time, in seconds.

Range: 120 through 1800 seconds

Default: 140 seconds

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Recovery Time and Maximum Recovery Time | 962](#)

session (Protocols LDP)

Syntax

```
session session-address {
  authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
  authentication-key authentication-key;
  authentication-key-chain authentication-key-chain;
  downstream-on-demand downstream-on-demand;
  (mtu-discovery | no-mtu-discovery);
  transport-address transport-address;
}
```

Hierarchy Level

```
[edit logical-systems name protocols ldp],
[edit logical-systems name routing-instances name protocols ldp],
[edit protocols ldp],
[edit routing-instances name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

authentication-algorithm statement introduced in Junos OS Release 7.6.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

transport-address option introduced in Junos OS Release 19.1R1 for all platforms.

Description

Configure LDP session parameters by specifying the session destination address.

Options

session-address—Session destination address.

authentication-algorithm—Authentication algorithm name.

Values:

- **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.

authentication-key—MD5 authentication key.

authentication-key-chain—Key chain name.

downstream-on-demand—Configure downstream on demand label distribution mode.

mtu-discovery—Enable TCP path MTU discovery.

no-mtu-discovery—Disable TCP path MTU discovery.

transport-address *transport-address*—IP address used for TCP sessions to the LDP neighbors that have targeted-LDP adjacencies.

You cannot configure this statement for LDP sessions that have discovered adjacencies.

The transport-address configuration can be rejected at the time of validation, if there is no interface with the configured IP address.

Required Privilege Level

routing

RELATED DOCUMENTATION

| [Configuring the TCP MD5 Signature for LDP Sessions](#) | **1118**

session-group

Syntax

```
session-group name {
  authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
  authentication-key authentication-key;
  authentication-key-chain authentication-key-chain;
  downstream-on-demand;
  (mtu-discovery | no-mtu-discovery);
  transport-address transport-address;
}
```

Hierarchy Level

```
[edit logical-systems name protocols ldp],
[edit logical-systems name routing-instances name protocols ldp],
[edit protocols ldp],
[edit routing-instances name protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 16.1.

transport-address option introduced in Junos OS Release 19.1R1 for all platforms.

Description

Specify the prefix address of the aggregated group of LDP neighbors for the remote end of the LDP session.

The **session-group** statement is useful when an LDP neighbor is dynamic; for instance, in the case of remote loop free alternate (LFA), a targeted or indirect LDP neighbor is automatically picked from any of the nodes in the network.

The **session-group** statement can also be used for configuring authentication for a prefix group, as LDP authentication for all sessions or authentication at the interface level is not supported.

The remaining statements are explained separately. See [CLI Explorer](#).

Options

name—Session destination address or prefix length.

authentication-algorithm—Authentication algorithm name.

Values:

- *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

authentication-key—MD5 authentication key.

authentication-key-chain—Authentication key chain name.

downstream-on-demand—Configure downstream on demand label distribution mode.

mtu-discovery | **no-mtu-discovery**—Enable and disable TCP path MTU discovery, respectively.

transport-address *transport-address*—IP address used for TCP sessions to the LDP neighbors that have targeted-LDP adjacencies, and fall under the same IP subnet.

You cannot configure this statement for LDP sessions that have discovered adjacencies.

The transport-address configuration can be rejected at the time of validation, if there is no interface with the configured IP address.

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](#) | 1118

session-protection

Syntax

```
session-protection {  
    timeout seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Description

Configure when an LDP session is torn down and resignaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.

Options

timeout *seconds*—Time in seconds before the LDP session is torn down and resignaled.

Range: 1 through 65,535 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LDP Session Protection](#) | 1120

source-packet-routing

Syntax

```
mapping-server-entry mapping-server-entry;
```

Hierarchy Level

```
[edit logical-systems name routing-instances name routing-options],  
[edit logical-systems name routing-options],  
[edit routing-instances name routing-options],  
[edit routing-options]
```

Release Information

Statement introduced in Junos OS Release 18.2R1.

Description

Configure interoperability between islands of devices supporting only segment routing and LDP in an LDP network domain where there is gradual deployment of segment routing.

Required Privilege Level

routing

stream-protection (Multicast-Only Fast Reroute)

Syntax

```
stream-protection {
  mofrr-asm-starg;
  mofrr-disjoint-upstream-only;
  mofrr-no-backup-join;
  mofrr-primary-path-selection-by-routing;
  policy policy-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name routing-options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Release 17.4R1 for QFX Series switches.

Description

Enable multicast-only fast reroute (MoFRR) on a routing or switching device. MoFRR minimizes packet loss in a network when there is a link failure.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Understanding Multicast-Only Fast Reroute | 1014](#)

Understanding Multicast-Only Fast Reroute on Switches

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain

Example: Configuring Multicast-Only Fast Reroute in a PIM Domain on Switches

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

strict-targeted-hellos

Syntax

```
strict-targeted-hellos;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Enabling Strict Targeted Hello Messages for LDP | 937](#)

targeted-hello

Syntax

```
targeted-hello {  
    hello-interval seconds;  
    hold-time seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced in Junos OS Release 9.5.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify the LDP timer and LDP hold time for targeted hellos.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the LDP Timer for Hello Messages | 934](#)

[Configuring the Delay Before LDP Neighbors Are Considered Down | 935](#)

traceoptions (Protocols LDP)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

match-on address option for the **filter** flag modifier added in Junos OS Release 10.4.

nsr-synchronization and **p2mp-nsr-synchronization** operations for **flag** statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Specify LDP protocol-level trace options.

Default

The default LDP protocol-level trace options are inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **ldp-log**. We recommend that you place LDP tracing output in the file **ldp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **address**—Operation of address and address withdrawal messages
- **binding**—Label-binding operations
- **error**—Error conditions
- **event**—Protocol events
- **initialization**—Operation of initialization messages
- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **nsr-synchronization**— Nonstop active routing synchronization events
- **p2mp-nsr-synchronization**—Point-to-multipoint nonstop active routing synchronization events
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags (see also the **filter** flag modifier)
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
 - **match-on**—Match on argument specified. The **match-on** option has the following suboptions:
 - **address**—Filter based on the source and destination addresses of packets. Available for the **packets** flag option only.
 - **fec**—Filter based on the FEC associated with the traced object.
 - **policy *policy-name***—Specify the filter policy.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent all users from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Tracing LDP Protocol Traffic | 1127](#)

Network Management and Monitoring Guide

track-igp-metric

Syntax

```
track-igp-metric;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],  
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],  
[edit protocols ldp],  
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Statement introduced in Junos OS Release 18.4R1 on MX5, MX10, MX40, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, and MX10008 routers.

Statement introduced in Junos OS Release 18.4R1 on PTX1000, PTX3000, PTX5000, PTX10002-60C, PTX10008, and PTX10016 routers.

Description

Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LDP to Use the IGP Route Metric](#) | 1116

track-igp-metric (LSP)

Syntax

```
track-igp-metric <install-v4-prefixes> <install-v6-prefixes>;
```

Hierarchy Level

The hierarchy level for **track-igp-metric** globally enabled for all LSPs:

```
[edit protocols mpls]
```

The hierarchy level for **track-igp-metric** at the per LSP level:

```
[edit protocols mpls label-switched-pathpathname],
```

Release Information

Statement introduced in Junos OS Release 18.4R1.

Description

Track IGP metric for LSP install prefixes

Options

install-v4-prefixes—Track IGP metric for IPV4 prefixes.

install-v6-prefixes—Track IGP metric for IPV6 prefixes.

Required Privilege Level

routing

RELATED DOCUMENTATION

[Install Prefix IGP Overview](#) | 26

traffic-statistics (Protocols LDP)

Syntax

```
traffic-statistics {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-penultimate-hop;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit routing-instances routing-instance-name protocols ldp]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.

Options

file *filename*—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 **/var/log**.

files *number*—(Optional) Maximum number of LDP statistics files. When a statistics file named ***ldp-stat*** reaches its maximum size, it is renamed ***ldp-stat.0***, then ***ldp-stat.1***, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you also must include the **size** statement to specify the maximum file size.

interval *seconds*—(Optional) Specify the interval at which the statistics are polled and written to the file.

Default: 300 seconds (5 minutes)

no-penultimate-hop—(Optional) Do not collect traffic statistics on the penultimate hop router.

no-world-readable—(Optional) Prevent all users from reading the log file.

size *size*—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named *ldp-stat* reaches this size, it is renamed *ldp-stat.0*. When *ldp-stat* again reaches this size, *ldp-stat.0* is renamed *ldp-stat.1* and *ldp-stat* is renamed *ldp-stat.0*. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics 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 also must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable log file access for all users.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Collecting LDP Statistics](#) | 1124

transport-address

Syntax

```
transport-address (address | interface | router-id);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp],
[edit logical-systems logical-system-name protocols ldp interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ldp],
[edit protocols ldp],
[edit protocols ldp interface interface-name],
[edit routing-instances routing-instance-name protocols ldp interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

address option introduced in Junos OS Release 19.1R1 for all platforms.

Description

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.

Default

router-id

Options

address—Use the specified IP address as the transport address for the LDP session.

interface—Use the first IP address on the interface as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

router-id—Use router identifier as the transport address. Unless otherwise configured, the router identifier is the loopback address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Specifying the Transport Address Used by LDP](#) | 968

version (BFD)

Syntax

```
version (0 | 1 | automatic);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ldp oam bfd-liveness-detection],
[edit logical-systems logical-system-name protocols ldp oam fec address bfd-liveness-detection],
[edit system services dhcp-local-server liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],
[edit forwarding-options dhcp-relay liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],
[edit system services dhcp-local-server group group-name liveness-detection method bfd],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay group group-name liveness-detection method bfd],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method bfd],
[edit protocols ldp oam bfd-liveness-detection],
[edit protocols ldp oam fec address bfd-liveness-detection]
```

Release Information

Statement introduced in Junos OS Release 12.1.

Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description

Configure the BFD protocol version to detect.

Options

0—Use BFD protocol version 0.

1—Use BFD protocol version 1.

automatic—Autodetect the BFD protocol version.

Default: automatic

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Group Liveness Detection with BFD for DHCP Local Server Clients

Example: Configuring Global Liveness Detection with BFD for DHCP Relay Agent Clients

[Configuring BFD for LDP LSPs](#) | **975**

CCC and TCC Configuration Statements

IN THIS CHAPTER

- connections (Circuits) | 3000
- encapsulation (Logical Interface) | 3002
- encapsulation | 3007
- interface-switch | 3014
- l2circuit-control-passthrough | 3015
- lsp-switch | 3016
- output-interface (CCC) | 3017
- p2mp-receive-switch | 3018
- p2mp-transmit-switch | 3019
- remote-interface-switch | 3020

connections (Circuits)

Syntax

```
connections {
  interface-switch connection-name {
    interface interface-name.unit-number;
  }
  lsp-switch connection-name {
    transmit-lsp label-switched-path;
    receive-lsp label-switched-path;
  }
  p2mp-receive-switch {
    output-interface [ interface-name.unit-number ];
    receive-p2mp-lsp receiving-point-to-multipoint-lsp;
  }
  p2mp-transmit-switch {
    input-interface interface-name.unit-number;
    transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;
  }
  remote-interface-switch connection-name {
    interface interface-name.unit-number;
    receive-lsp label-switched-path;
    transmit-lsp label-switched-path;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define the connection between two circuits in a CCC connection.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

NOTE: The **edit logical-systems** hierarchy is not available on QFabric systems.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Layer 2 Switching Cross-Connects Using CCC | 1655](#)

[Configuring MPLS LSP Tunnel Cross-Connects Using CCC | 1665](#)

[Configuring TCC | 1671](#)

[Configuring CCC Switching for Point-to-Multipoint LSPs | 1681](#)

encapsulation (Logical Interface)

Syntax

```
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-mlppp-llc | atm-nlpid | atm-ppp-llc |
atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-tcc-vc-mux | atm-vc-mux | ether-over-atm-llc |
ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet | ethernet-ccc | ethernet-vpls |
ethernet-vpls-fr | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-ppp |
frame-relay-tcc | gre-fragmentation | multilink-frame-relay-end-to-end | multilink-ppp | ppp-over-ether |
ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc | vlan-tcc | vlan-vpls | vxlan);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number],
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number],
[edit interfaces rlsq number unit logical-unit-number]
[edit protocols evpn]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers

(**ethernet**, **vlan-ccc**, and **vlan-tcc** options only).

Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Metro Routers. Only the **atm-ccc-cell-relay** and **atm-ccc-vc-mux** options are supported on ACX Series routers.

Statement introduced in Junos OS Release 17.3R1 for QFX10000 Series switches (**ethernet-ccc** and **vlan-ccc** options only).

Description

Configure a logical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.

Options

atm-ccc-cell-relay—Use ATM cell-relay encapsulation.

atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

atm-cisco-nlpid—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

atm-ppp-llc—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.

atm-ppp-vc-mux—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.

atm-snap—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.

atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.

NOTE: The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.

ether-vpls-over-ppp—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE router over a time-division multiplexing (TDM) link. This encapsulation type enables the PE router to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—Use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface on M120 and M320 routers with Intelligent Queuing 2 (IQ2) PICs, and on MX Series routers with MPCs.

ppp-over-ether-over-atm-llc—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

vxlan—Use VXLAN data plane encapsulation for EVPN.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Layer 2 Switching Cross-Connects Using CCC | 1655](#)

[Configuring the Encapsulation for Layer 2 Switching TCCs | 1671](#)

[Configuring Interface Encapsulation on Logical Interfaces](#)

[Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects | 1667](#)

[Circuit and Translational Cross-Connects Overview](#)

[Identifying the Access Concentrator](#)

[Configuring ATM Interface Encapsulation](#)

[Configuring VLAN and Extended VLAN Encapsulation](#)

[Configuring ATM-to-Ethernet Interworking](#)

[Configuring Interface Encapsulation on PTX Series Packet Transport Routers](#)

[Configuring CCC Encapsulation for Layer 2 VPNs](#)

[Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits](#)

[Configuring ATM for Subscriber Access](#)

[Understanding CoS on ATM IMA Pseudowire Interfaces Overview](#)

[Configuring Policing on an ATM IMA Pseudowire](#)

encapsulation

List of Syntax

[Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series on page 3007](#)

[Syntax for Physical Interfaces: SRX Series on page 3007](#)

[Syntax for Logical Interfaces: SRX Series on page 3007](#)

Syntax for Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

```
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc | ethernet-bridge | ethernet-ccc
| ethernet-over-atm | ethernet-tcc | ethernet-vpls | ethernet-vpls-fr | ether-vpls-over-atm-llc | ethernet-vpls-ppp
| extended-frame-relay-ccc | extended-frame-relay-ether-type-tcc | extended-frame-relay-tcc |
extended-vlan-bridge | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services |
flexible-frame-relay | frame-relay | frame-relay-ccc | frame-relay-ether-type | frame-relay-ether-type-tcc |
frame-relay-port-ccc | frame-relay-tcc | generic-services | multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc
| vlan-ccc | vlan-vci-ccc | vlan-vpls);
```

Syntax for Physical Interfaces: SRX Series

```
encapsulation (ether-vpls-ppp | ethernet-bridge | ethernet-ccc | ethernet-tcc | ethernet-vpls |
extended-frame-relay-ccc | extended-frame-relay-tcc | extended-vlan-bridge | extended-vlan-ccc |
extended-vlan-tcc | extended-vlan-vpls | flexible-ethernet-services | frame-relay-port-ccc | vlan-ccc | vlan-vpls);
```

Syntax for Logical Interfaces: SRX Series

```
encapsulation ( dix | ether-vpls-fr | frame-relay-ppp | ppp-over-ether | vlan-bridge | vlan-ccc | vlan-tcc | vlan-vpls );
```

Physical Interfaces: M Series, MX Series, QFX Series, T Series, PTX Series

```
[edit interfaces interface-name],
[edit interfaces rlsq number:number]
```

Logical Interfaces

```
[edit interfaces interface-name unit logical-unit-number ]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5.


Statement introduced in Junos OS Release 11.1 for EX Series switches.

Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (**flexible-ethernet-services**, **ethernet-ccc**, and **ethernet-tcc** options only).

Description

For M Series, MX Series, QFX Series, T Series, PTX Series, specify the physical link-layer encapsulation type.

For SRX Series, specify logical link layer encapsulation.



NOTE: Not all encapsulation types are supported on the switches. See the switch CLI.

Default

ppp—Use serial PPP encapsulation.

Physical Interface Options and Logical Interface Options

[Warning: element unresolved in stylesheets: <title> (in <config-options>). This is probably a new element that is not yet supported in the stylesheets.]

Physical Interface Options and Logical Interface Options

For physical interfaces:

NOTE: Frame Relay, ATM, PPP, SONET, and SATSOP options are not supported on EX Series switches.

- **atm-ccc-cell-relay**—Use ATM cell-relay encapsulation.
- **atm-pvc**—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).
- **cisco-hdlc**—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:
 - CCC version (**cisco-hdlc-ccc**)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
 - TCC version (**cisco-hdlc-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **cisco-hdlc-ccc**—Use Cisco-compatible HDLC framing on CCC circuits.
- **cisco-hdlc-tcc**—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.
- **ethernet-bridge**—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.
- **ethernet-over-atm**—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.
- **ethernet-tcc**—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

- **ethernet-vpls**—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.
- **ethernet-vpls-fr**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.
- **ethernet-vpls-ppp**—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.
- **ether-vpls-over-atm-llc**—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.
- **extended-frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. When you use this encapsulation type, you can configure the **ccc** family only.
- **extended-frame-relay-ether-type-tcc**—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.
- **extended-frame-relay-tcc**—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.
- **extended-vlan-bridge**—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.
- **extended-vlan-ccc**—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.
- **extended-vlan-tcc**—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

- **extended-vlan-vpls**—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

- **flexible-ethernet-services**—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and for Gigabit Ethernet interfaces, use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.
- **flexible-frame-relay**—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.
- **frame-relay**—Use Frame Relay encapsulation is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation.
- **frame-relay-ccc**—Use Frame Relay encapsulation on CCC circuits. This encapsulation is same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- **frame-relay-ether-type**—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload.

NOTE: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

- **frame-relay-ether-type-tcc**—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. This encapsulation is Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC.

- **frame-relay-port-ccc**—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. When you use this encapsulation type, you can configure the **ccc** family only.
- **frame-relay-tcc**—This encapsulation is similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **generic-services**—Use generic services encapsulation for services with a hierarchical scheduler.
- **multilink-frame-relay-uni-nni**—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.
-
- **ppp**—Use serial PPP encapsulation. This encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation.
- **ppp-ccc**—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.
- **ppp-tcc**—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.
- **vlan-ccc**—Use Ethernet VLAN encapsulation on CCC circuits. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.
- **vlan-vci-ccc**—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.
- **vlan-vpls**—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

NOTE:

- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
- Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure **family inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.

For logical interfaces:

- **frame-relay**—Configure a Frame Relay encapsulation when the physical interface has multiple logical units, and the units are either point to point or multipoint.
- **multilink-frame-relay-uni-nni**—Link services interfaces functioning as FRF.16 bundles can use Multilink Frame Relay UNI NNI encapsulation.
- **ppp**—For normal mode (when the device is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface.
- **ppp-over-ether**—This encapsulation is used for underlying interfaces of pp0 interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding Physical Encapsulation on an Interface

Configuring Interface Encapsulation on Physical Interfaces

Configuring CCC Encapsulation for Layer 2 VPNs

[Configuring Layer 2 Switching Cross-Connects Using CCC | 1655](#)

Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits

Configuring ATM Interface Encapsulation

Configuring ATM-to-Ethernet Interworking

Configuring VLAN and Extended VLAN Encapsulation

Configuring VLAN and Extended VLAN Encapsulation

Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces

Configuring Interfaces for Layer 2 Circuits

Configuring Interface Encapsulation on PTX Series Packet Transport Routers

[Configuring MPLS LSP Tunnel Cross-Connects Using CCC | 1665](#)

[Configuring TCC | 1671](#)

Configuring VPLS Interface Encapsulation

Configuring Interfaces for VPLS Routing

Defining the Encapsulation for Switching Cross-Connects

[Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) | 1460](#)

interface-switch

Syntax

```
interface-switch connection-name {
    interface interface-name.unit-number;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],
[edit protocols connections]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure Layer 2 switching cross-connects. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.

For Layer 2 switching cross-connects to work, you must also configure MPLS.

Options

connection-name—Connection name (up to 128 characters in Junos 12.3 and later).

interface interface-name.unit-number—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration

RELATED DOCUMENTATION

[Configuring the CCC Connection for Layer 2 Switching Cross-Connects | 1661](#)

[Defining the Connection for Switching Cross-Connects](#)

[MPLS Applications User Guide](#)

l2circuit-control-passthrough

Syntax

```
l2circuit-control-passthrough;
```

Hierarchy Level

```
[edit forwarding-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R1 for PTX Series routers.

Description

Configure the device to allow LACP, LLDP, OAM LFM, and OAM CFM packets to cross the Layer 2 circuit. If the **l2circuit-control-passthrough** statement is not configured, LACP, LLDP, OAM LFM, and OAM CFM packets are classified as control packets and are not transmitted across the Layer 2 circuit.

NOTE: For MX Series routers, the functionality that the **l2circuit-control-passthrough** command provides is performed automatically.

Default

By default, this statement is not configured.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

forwarding-options

[Configuring Layer 2 Switching Cross-Connects Using CCC | 1655](#)

lsp-switch

Syntax

```
lsp-switch connection-name {  
    transmit-lsp label-switched-path;  
    receive-lsp label-switched-path;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],  
[edit protocols connections]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure Layer 2 switching cross-connects.

Options

connection-name—Connection name.

receive-lsp label-switched-path—Name of the LSP from the connection's source.

transmit-lsp label-switched-path—Name of the LSP to the connection's destination.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Connection for Layer 2 Switching TCCs](#) | 1675

output-interface (CCC)

Syntax

```
output-interface [interface-name 1 interface-name n];
```

Hierarchy Level

```
[edit protocols connections p2mp-transmit-switch p2mp-transmit-switch-name]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Description

Specify one or more output interfaces to switch traffic on an incoming CCC interface to one or more outgoing CCC interfaces.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring CCC Switching for Point-to-Multipoint LSPs](#) | 1681

p2mp-receive-switch

Syntax

```
p2mp-receive-switch point-to-multipoint-switch-name {  
    output-interface [ interface-name.unit-number ];  
    receive-p2mp-lsp receiving-point-to-multipoint-lsp;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],  
[edit protocols connections]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the CCC switch for a point-to-multipoint LSP on the egress PE router.

Options

point-to-multipoint-switch-name—Point-to-multipoint CCC receive switch name.

output-interface *interface-name.unit-number*—Name of the egress interfaces for the point-to-multipoint LSP traffic. You can configure multiple output interfaces.

receive-p2mp-lsp *receiving-point-to-multipoint-lsp*—Name of the point-to-multipoint LSP that is switched to the output interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers](#) | 1683

p2mp-transmit-switch

Syntax

```
p2mp-transmit-switch point-to-multipoint-transmit-switch-name {  
    input-interface interface-name.unit-number;  
    transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],  
[edit protocols connections]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the CCC switch for the point-to-multipoint LSP on the ingress PE router.

Options

point-to-multipoint-transmit-switch-name—Point-to-multipoint CCC transmit switch name.

input-interface input-interface-name.unit-number—Specify the name of the interface carrying incoming traffic to be switched to the point-to-multipoint LSP.

transmit-p2mp-lsp transmitting-point-to-multipoint-lsp—Specify the name of the point-to-multipoint LSP carrying traffic to the CCC switch on the egress PE router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers](#) | 1682

remote-interface-switch

Syntax

```
remote-interface-switch connection-name {  
    interface interface-name.unit-number;  
    transmit-lsp label-switched-path;  
    receive-lsp label-switched-path;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols connections],  
[edit protocols connections]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure MPLS LSP tunnel cross-connects.

Options

connection-name—Connection name.

interface interface-name.unit-number—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.

receive-lsp label-switched-path—Name of the LSP from the connection's source.

transmit-lsp label-switched-path—Name of the LSP to the connection's destination.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS LSP Tunnel Cross-Connects Using CCC](#) | 1665

GMPLS Configuration Statements

IN THIS CHAPTER

- address (Peer) | 3022
- control-channel (Protocols Link Management Peer) | 3023
- dead-interval | 3024
- disable (GMPLS) | 3025
- disable (OSPF) | 3026
- export (Protocols BGP) | 3028
- hello-dead-interval | 3029
- hello-interval (LMP) | 3030
- hello-interval (Protocols OSPF) | 3031
- import | 3033
- instance-type | 3035
- interface (Protocols Link Management) | 3038
- label-switched-path (Protocols Link Management) | 3039
- link-management | 3040
- lmp-control-channel | 3042
- lmp-protocol | 3043
- local-address (Protocols Link Management) | 3044
- l2circuit | 3045
- passive (Protocols Link Management) | 3047
- peer (Protocols LMP) | 3048
- peer-interface (Protocols OSPF) | 3049
- remote-address (for LMP Control Channel) | 3050
- remote-address (for LMP Traffic Engineering) | 3051
- remote-id | 3052
- retransmission-interval | 3053
- retransmit-interval (OSPF) | 3054
- retry-limit (Protocols Link Management) | 3056
- route-distinguisher | 3057

- [te-link | 3061](#)
- [traceoptions \(Protocols Link Management\) | 3063](#)
- [transit-delay \(OSPF\) | 3065](#)
- [upstream-label | 3067](#)
- [vrf-target | 3068](#)

address (Peer)

Syntax

```
address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],  
[edit protocols link-management peer peer-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the ID of the peer.

Default

The loopback address is advertised.

Options

ip-address—IP address of the peer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the ID for LMP Peers | 1364](#)

control-channel (Protocols Link Management Peer)

Syntax

```
control-channel control-channel-interface;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],  
[edit protocols link-management peer peer-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the control channel interface for the peer.

Options

control-channel-interface—Name of the control channel interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring LMP Peers](#) | 1363

dead-interval

Syntax

```
dead-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id peer-interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id virtual-link],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area
  area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) area area-id
  interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) area area-id
  virtual-link],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast
  | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit protocols ospf area area-id peer-interface interface-name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols (ospf | ospf3) area area-id virtual-link],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id virtual-link],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)
  area area-id interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the **realm** statement introduced in Junos OS Release 9.2.

Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.

Description

Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor.

Options

seconds—Interval to wait.

Range: 1 through 65,535 seconds

Default: Four times the hello interval—40 seconds (broadcast and point-to-point networks); 120 seconds (nonbroadcast multiple access (NBMA) networks)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring OSPF Timers

[Configuring RSVP and OSPF for LMP Peer Interfaces](#) | 1369

[hello-interval](#) | 3031

disable (GMPLS)**Syntax**

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],  
[edit protocols link-management te-link te-link-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Disable a traffic engineering link.

Default

The configured object is enabled (operational) unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Disabling the Traffic Engineering Link for LMP Peers](#) | 1368

disable (OSPF)

Syntax

```
disable;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name protocols ospf area area-id peer-interfaceinterface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) virtual-link],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) virtual-link],
[edit logical-systems logical-system-name routing-instances routing-instances protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit protocols (ospf | ospf3)],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols (ospf | ospf3) virtual-link],
[edit protocols ospf area area-id peer-interface interface-name],
[edit protocols ospf area area-id virtual-link neighbor-id router-id transit-area area-id],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3)],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) virtual-link],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the **realm** statement introduced in Junos OS Release 9.2.

Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Disable OSPF, an OSPF interface, or an OSPF virtual link.

By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id* topology *name*]** hierarchy level.

NOTE: If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

Default

The configured object is enabled (operational) unless explicitly disabled.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Understanding OSPF Configurations

[Configuring RSVP and OSPF for LMP Peer Interfaces](#) | 1369

export (Protocols BGP)

Syntax

```
export [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name
  neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply one or more policies to routes being exported from the routing table into BGP.

If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no routes match the filters, the routing table exports into BGP only the routes that it learned from BGP. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of **metric 500** defined in the next policy.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Routing Policies to Control BGP Route Advertisements](#)

[Routing Policies, Firewall Filters, and Traffic Policers User Guide](#)

[import](#) | [3033](#)

hello-dead-interval

Syntax

```
hello-dead-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],  
[edit protocols link-management peer peer-name lmp-protocol]
```

Release Information

Statement introduced in Junos OS Release 8.0.

Description

Specify how long the Link Management Protocol (LMP) waits before declaring the control channel to be dead. This is an interval during which the router receives no LMP hello packets from the neighbor on a control that is active or up.

Options

milliseconds—Interval to wait before declaring the control channel to be dead.

Range: 500 through 300,000

Default: 500 milliseconds (three times the hello interval)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Hello Message Intervals for LMP Control Channels](#) | [1366](#)

[hello-interval \(LMP\)](#) | [3030](#)

hello-interval (LMP)

Syntax

```
hello-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],  
[edit protocols link-management peer peer-name lmp-protocol]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specify how often the router sends Link Management Protocol (LMP) hello packets.

Options

milliseconds—Length of time between hello packets.

Range: 150 through 300,000

Default: 150 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Hello Message Intervals for LMP Control Channels | 1366](#)

[hello-dead-interval | 3029](#)

hello-interval (Protocols OSPF)

Syntax

```
hello-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id peer-interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id virtual-link],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area
  area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) area area-id
  interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) area area-id
  virtual-link],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast
  | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit protocols ospf area area-id peer-interface interface-name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols (ospf | ospf3) area area-id virtual-link],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id virtual-link],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)
  area area-id interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the **realm** statement introduced in Junos OS Release 9.2.

Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.

Description

Specify how often the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network.

Options

seconds—Time between hello packets, in seconds.

Range: 1 through 255 seconds

Default: 10 seconds (broadcast and point-to-point networks); 30 seconds (nonbroadcast multiple access [NBMA] networks)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring OSPF Timers

[Configuring RSVP and OSPF for LMP Peer Interfaces](#) | **1369**

[dead-interval](#) | **3024**

import

Syntax

```
import [ policy-names ];
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols bgp],
[edit logical-systems logical-system-name protocols bgp group group-name],
[edit logical-systems logical-system-name protocols bgp group group-name neighbor address],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols bgp group group-name
  neighbor address],
[edit protocols bgp],
[edit protocols bgp group group-name],
[edit protocols bgp group group-name neighbor address],
[edit routing-instances routing-instance-name protocols bgp],
[edit routing-instances routing-instance-name protocols bgp group group-name],
[edit routing-instances routing-instance-name protocols bgp group group-name neighbor address]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.

If you specify more than one policy, they are evaluated in the order specified, from left to right, and the first matching filter is applied to the route. If no match is found, BGP places into the routing table only those routes that were learned from BGP routing devices. The policy framework software evaluates the routing policies in a chain sequentially. If an action specified in one of the policies manipulates a route characteristic, the policy framework software carries the new route characteristic forward during the evaluation of the remaining policies. For example, if the action specified in the first policy of a chain sets a route's metric to 500, this route matches the criterion of **metric 500** defined in the next policy.

It is also important to understand that in Junos OS, although an import policy (inbound route filter) might reject a route, not use it for traffic forwarding, and not include it in an advertisement to other peers, the router retains these routes as hidden routes. These hidden routes are not available for policy or routing purposes. However, they do occupy memory space on the router. A service provider filtering routes to

control the amount of information being kept in memory and processed by a router might want the router to entirely drop the routes being rejected by the import policy.

Hidden routes can be viewed by using the **show route receive-protocol bgp neighbor-address hidden** command. The hidden routes can then be retained or dropped from the routing table by configuring the **keep all | none** statement at the **[edit protocols bgp]** or **[edit protocols bgp group group-name]** hierarchy level.

The rules of BGP route retention are as follows:

- By default, all routes learned from BGP are retained, except those where the AS path is looped. (The AS path includes the local AS.)
- By configuring the **keep all** statement, all routes learned from BGP are retained, even those with the local AS in the AS path.
- By configuring the **keep none** statement, all routes received are discarded. When this statement is configured and the inbound policy changes, Junos OS re-advertises all the routes advertised by the peer.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring BGP Interactions with IGPs

Configuring Routing Policies to Control BGP Route Advertisements

Understanding Routing Policies

[export](#) | [3028](#)

instance-type

Syntax

```
instance-type type;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],  
[edit routing-instances routing-instance-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

virtual-switch and **layer2-control** options introduced in Junos OS Release 8.4.

Statement introduced in Junos OS Release 9.2 for EX Series switches.

Statement introduced in Junos OS Release 11.3 for the QFX Series.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

mpls-internet-multicast option introduced in Junos OS Release 11.1 for the EX Series, M Series, MX Series, and T Series.

evpn option introduced in Junos OS Release 13.2 for MX 3D Series routers.

evpn option introduced in Junos OS Release 17.3 for the QFX Series.

forwarding option introduced in Junos OS Release 14.2 for the PTX Series.

mpls-forwarding option introduced in Junos OS Release 16.1 for the MX Series.

evpn-vpws option introduced in Junos OS Release 17.1 for MX Series routers.

Support for logical systems on MX Series routers added in Junos OS Release 17.4R1.

Description

Define the type of routing instance.

Options

NOTE: On ACX Series routers, you can configure only the forwarding, virtual router, and VRF routing instances.

type—Can be one of the following:

- **evpn**—(MX 3D Series routers, QFX switches and EX9200 switches)—Enable an Ethernet VPN (EVPN) on the routing instance.
hierarchy level.
- **evpn-vpws**—Enable an Ethernet VPN (EVPN) Virtual Private Wire Service (VPWS) on the routing instance.

- **forwarding**—Provide support for filter-based forwarding, where interfaces are not associated with instances. All interfaces belong to the default instance. Other instances are used for populating RPD learned routes. For this instance type, there is no one-to-one mapping between an interface and a routing instance. All interfaces belong to the default instance inet.0.
- **l2backhaul-vpn**—Provide support for Layer 2 wholesale VLAN packets with no existing corresponding logical interface. When using this instance, the router learns both the outer tag and inner tag of the incoming packets, when the **instance-role** statement is defined as **access**, or the outer VLAN tag only, when the **instance-role** statement is defined as **nni**.
- **l2vpn**—Enable a Layer 2 VPN on the routing instance. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **layer2-control**—(MX Series routers only) Provide support for RSTP or MSTP in customer edge interfaces of a VPLS routing instance. This instance type cannot be used if the customer edge interface is multihomed to two provider edge interfaces. If the customer edge interface is multihomed to two provider edge interfaces, use the default BPDU tunneling.
- **mpls-forwarding**—(MX Series routers only) Allow filtering and translation of route distinguisher (RD) values in IPv4 and IPv6 VPN address families on both routes received and routes sent for selected BGP sessions. In particular, for Inter-AS VPN Option-B networks, this option can prevent the malicious injection of VPN labels from one peer AS boundary router to another.
- **mpls-internet-multicast**—(EX Series, M Series, MX Series, and T Series routers only) Provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.
- **no-forwarding**—This is the default routing instance. Do not create a corresponding forwarding instance. Use this routing instance type when a separation of routing table information is required. There is no corresponding forwarding table. All routes are installed into the default forwarding table. IS-IS instances are strictly nonforwarding instance types.
- **virtual-router**—Enable a virtual router routing instance. This instance type is similar to a VPN routing and forwarding instance type, but used for non-VPN-related applications. You must configure the **interface** statement for this type of routing instance. You do not need to configure the **route-distinguisher**, **vrf-import**, and **vrf-export** statements.
- **virtual-switch**—(MX Series routers, EX9200 switches, and QFX switches only) Provide support for Layer 2 bridging. Use this routing instance type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.

- **vpls**—Enable VPLS on the routing instance. Use this routing instance type for point-to-multipoint LAN implementations between a set of sites in a VPN. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.
- **vrf**—VPN routing and forwarding (VRF) instance. Provides support for Layer 3 VPNs, where interface routes for each instance go into the corresponding forwarding table only. Required to create a Layer 3 VPN. Create a VRF table (**instance-name.inet.0**) that contains the routes originating from and destined for a particular Layer 3 VPN. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. You must configure the **interface**, **route-distinguisher**, **vrf-import**, and **vrf-export** statements for this type of routing instance.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring the Instance Type

Configuring EVPN Routing Instances

Configuring EVPN Routing Instances on EX9200 Switches

Configuring Virtual Router Routing Instances

Example: Configuring Filter-Based Forwarding on the Source Address

Example: Configuring Filter-Based Forwarding on Logical Systems

interface (Protocols Link Management)

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],  
[edit protocols link-management te-link te-link-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the egress router interface.

Options

interface-name—Name of the interface to the egress router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[LMP Configuration Overview](#) | 1359

label-switched-path (Protocols Link Management)

Syntax

```
label-switched-path lsp-name;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],  
[edit protocols link-management te-link te-link-name]
```

Release Information

Statement introduced in Junos OS Release 7.4.

Description

Specify the label-switched path (LSP) to be used by the forwarding adjacency.

Options

lsp-name—Name of the LSP.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Configuring Forwarding Adjacencies*

link-management

Syntax

```

link-management {
  peer peer-name {
    address ip-address;
    control-channel control-channel-interface;
    imp-control-channel control-channel-interface {
      remote-address ip-address;
    }
    imp-protocol {
      hello-dead-interval milliseconds;
      hello-interval milliseconds;
      passive;
      retransmission-interval milliseconds;
      retry-limit number;
    }
    te-link te-link-name;
  }
  te-link te-link-name {
    disable;
    interface interface-name {
      disable;
      local-address ip-address;
      remote-address ip-address;
      remote-id id-number;
    }
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}

```

Hierarchy Level

```

[edit logical-systems logical-system-name protocols],
[edit protocols]

```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable Link Management Protocol (LMP) on the router.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [LMP Configuration Overview](#) | 1359

Imp-control-channel

Syntax

```
imp-control-channel control-channel-interface {  
    remote-address ip-address;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],  
[edit protocols link-management peer peer-name]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specify the Link Management Protocol (LMP) control channel interface for the peer.

Options

control-channel-interface—Name of the control channel interface.

The remaining statement is described separately in this chapter.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the LMP Control Channel Interface for the Peer](#) | 1365

Imp-protocol

Syntax

```
imp-protocol {  
    hello-dead-interval milliseconds;  
    hello-interval milliseconds;  
    passive;  
    retransmission-interval milliseconds;  
    retry-limit number;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name],  
[edit protocols link-management peer peer-name]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Configure attributes of Link Management Protocol (LMP) to establish and maintain the LMP control channel for the peer.

Options

The statements are described separately in this chapter.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LMP Peers](#) | 1363

local-address (Protocols Link Management)

Syntax

```
local-address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],  
[edit logical-systems logical-system-name protocols link-management te-link te-link-name interface interface-name],  
[edit protocols link-management te-link te-link-name],  
[edit protocols link-management te-link te-link-name interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the local IP address associated with the traffic engineering link. If you configure the local IP address, you must also configure the [remote-address](#) statement.

Options

local-address—Local IP address of the traffic engineering link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Local IP Address for Traffic Engineering Links | 1361](#)

[Configuring the Local IP Address for Forwarding Adjacencies](#)

[remote-address \(for LMP Traffic Engineering\) | 3051](#)

I2circuit

Syntax

```

I2circuit {
  auto-sensing{
    password password;
  }
  local-switching {
    interface interface-name {
      description text;
      end-interface {
        interface interface-name;
        protect-interface interface-name;
      }
      ignore-mtu-mismatch;
      protect-interface interface-name;
    }
  }
  neighbor address {
    interface interface-name {
      backup-neighbor address;
      bandwidth (bandwidth | ctnumber bandwidth);
      community community-name;
      connection-protection;
      (control-word | no-control-word);
      description text;
      egress-protection;
      encapsulation-type type;
      ignore-encapsulation-mismatch;
      ignore-mtu-mismatch;
      mtu mtu-number;
      protect-interface interface-name;
      pseudowire-status-tlv hot-standby-vc-on;
      psn-tunnel-endpoint address;
      virtual-circuit-id identifier;
    }
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}

```

Hierarchy Level


```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.

Description

Enables a Layer 2 circuit.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

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

passive (Protocols Link Management)

Syntax

```
passive;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],  
[edit protocols link-management peer peer-name lmp-protocol]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specify that the router not configure the Link Management Protocol (LMP) control channels but wait for the remote peer to configure the LMP control channels.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Preventing the Local Peer from Initiating LMP Negotiation](#) | 1367

peer (Protocols LMP)

Syntax

```
peer peer-name {  
    address ip-address;  
    control-channel control-channel-interface;  
    lmp-control-channel control-channel-interface;  
    lmp-protocol {  
        hello-dead-interval milliseconds;  
        hello-interval milliseconds;  
        passive;  
        retransmission-interval milliseconds;  
        retry-limit number;  
    }  
    te-link te-link-name;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management],  
[edit protocols link-management]
```

Release Information

Statement introduced before Junos OS Release 7.4.

lmp-protocol statement and substatements added in Junos OS Release 8.1.

Description

Configure a network peer.

Options

peer-name—Name of the network peer.

The remaining statements are described separately in this chapter.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring LMP Peers](#) | 1363

peer-interface (Protocols OSPF)

Syntax

```
peer-interface interface-name {  
    disable;  
    dead-interval seconds;  
    hello-interval seconds;  
    retransmit-interval seconds;  
    transit-delay seconds;  
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id],  
[edit protocols ospf area area-id]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure a peer interface.

Options

interface-name—Name of the peer interface. To configure all interfaces, you can specify **all**.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring OSPFv2 Peer interfaces

[Configuring RSVP and OSPF for LMP Peer Interfaces](#) | 1369

Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP

remote-address (for LMP Control Channel)

Syntax

```
remote-address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-control-channel  
  control-channel-interface],  
[edit protocols link-management peer peer-name lmp-control-channel control-channel-interface]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specify the remote IP address for the Link Management Protocol (LMP) control channel interface.

Options

ip-address—Remote IP address mapped to the LMP control channel interface.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Remote IP Address for LMP Control Channels](#) | 1366

remote-address (for LMP Traffic Engineering)

Syntax

```
remote-address ip-address;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],  
[edit logical-systems logical-system-name protocols link-management te-link te-link-name interface interface-name],  
[edit protocols link-management te-link te-link-name],  
[edit protocols link-management te-link te-link-name interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the remote IP address for the traffic engineering link. If you configure the remote IP address, you must also configure the [local-address](#) statement.

Options

ip-address—Remote IP address mapped to the traffic engineering link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring the Remote IP Address for Traffic Engineering Links](#) | 1362

[Configuring the Remote IP Address for Forwarding Adjacencies](#)

[local-address \(Protocols Link Management\)](#) | 3044

remote-id

Syntax

```
remote-id id-number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management te-link te-link-name],  
[edit logical-systems logical-system-name protocols link-management te-link te-link-name interface interface-name],  
[edit protocols link-management te-link te-link-name],  
[edit protocols link-management te-link te-link-name interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the ID assigned to a traffic engineering link or an interface (resource) on the peer node.

Options

id-number—ID number for the remote device.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring the Remote ID for Traffic Engineering Links](#) | 1362

retransmission-interval

Syntax

```
retransmission-interval milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],  
[edit protocols link-management peer peer-name lmp-protocol]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specify how often Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel.

Options

milliseconds—Length of time between Config messages.

Range: 500 through 300,000

Default: 500 milliseconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[retry-limit \(Protocols Link Management\) | 3056](#)

[Controlling Message Exchange for LMP Control Channels | 1367](#)

retransmit-interval (OSPF)

Syntax

```
retransmit-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id peer-interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id virtual-link],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area
  area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) area area-id
  interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols (ospf | ospf3) area area-id
  virtual-link],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast
  | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit protocols ospf area area-id peer-interface interface-name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols (ospf | ospf3) area area-id virtual-link],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols (ospf | ospf3) area area-id virtual-link],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)
  area area-id interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the **realm** statement introduced in Junos OS Release 9.2.

Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.

Description

Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors.

Options

seconds—Interval to wait.

Range: 1 through 65,535 seconds

Default: 5 seconds

NOTE: You must configure LSA retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because Junos OS delays LSA acknowledgments by up to 2 seconds.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring OSPF Timers

[Configuring RSVP and OSPF for LMP Peer Interfaces](#) | 1369

retry-limit (Protocols Link Management)

Syntax

```
retry-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management peer peer-name lmp-protocol],  
[edit protocols link-management peer peer-name lmp-protocol]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Specify how many times the Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel without receiving an appropriate acknowledgment before it logs a message and restarts the LMP control channel configuration process.

Options

number—Maximum number of times messages are sent without receiving an acknowledgment.

Range: 3 through 1000

Default: 3

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[retransmission-interval | 3053](#)

[Controlling Message Exchange for LMP Control Channels | 1367](#)

route-distinguisher

Syntax

```
route-distinguisher (as-number:id | ip-address:id);
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn mesh-group
  mesh-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls mesh-group
  mesh-group-name],
[edit routing-instances routing-instance-name],
[edit routing-instances routing-instance-name protocols l2vpn mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Support at **[edit routing-instances *routing-instance-name* protocols vpls mesh-group *mesh-group-name*]**

hierarchy level introduced in Junos OS Release 11.2.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Support at **[edit routing-instances *routing-instance-name* protocols l2vpn mesh-group *mesh-group-name*]**

hierarchy level introduced in Junos OS Release 13.2.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.

Statement introduced in cRPD Release 19.4R1.

Description

Specify an identifier attached to a route, enabling you to distinguish to which VPN or virtual private LAN service (VPLS) the route belongs. Each routing instance must have a unique route distinguisher (RD) associated with it. The RD is used to place bounds around a VPN so that the same IP address prefixes can be used in different VPNs without having them overlap. If the instance type is **vrf**, the **route-distinguisher** statement is required.

For Layer 2 VPNs and VPLS, if you configure the **l2vpn-use-bgp-rules** statement, you must configure a unique RD for each PE router participating in the routing instance.

For other types of VPNs, we recommend that you use a unique RD for each provider edge (PE) router participating in specific routing instance. Although you can use the same RD on all PE routers for the same VPN routing instance, if you use a unique RD, you can determine the customer edge (CE) router from which a route originated within the VPN.

For Layer 2 VPNs and VPLSs, if you configure mesh groups, the RD in each mesh group must be unique.



CAUTION: We strongly recommend that if you change an RD that has already been configured, or change the routing-instance type from **virtual-router** to **vrf**, make the change during a maintenance window, as follows:

1. Deactivate the routing instance.
2. Change the RD.
3. Activate the routing instance.

This is not required if you are configuring the RD for the first time.

Options

as-number:number—as-number is an assigned AS number, and **number** is any 2-byte or 4-byte value. The AS number can be from 1 through 4,294,967,295. If the AS number is a 2-byte value, the administrative number is a 4-byte value. If the AS number is 4-byte value, the administrative number is a 2-byte value. An RD consisting of a 4-byte AS number and a 2-byte administrative number is defined as a type 2 RD in RFC 4364 *BGP/MPLS IP VPNs*.

NOTE: In Junos OS Release 9.1 and later, the numeric range for AS numbers is extended to provide BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. All releases of Junos OS support 2-byte AS numbers. To configure an RD that includes a 4-byte AS number, append the letter “L” to the end of the AS number. For example, an RD with the 4-byte AS number 7,765,000 and an administrative number of 1,000 is represented as **77765000L:1000**.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in the plain-number format is represented as 1.10 in AS dot notation format.

ip-address:id—IP address (**ip-address** is a 4-byte value) within your assigned prefix range and a 2-byte value for the **id**. The IP address can be any globally unique unicast address.

Range: 0 through 4,294,967,295 ($2^{32} - 1$). If the router you are configuring is a BGP peer of a router that does not support 4-byte AS numbers, you need to configure a local AS number. For more information, see *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

NOTE: For Ethernet VPN (EVPN), an RD that includes zero as the **id** value is reserved for the default EVPN routing instance by default. Because the same RD cannot be assigned for two routing instances, using a **ip-address:id** RD for another routing instance (default-switch), where the **id** value is zero, throws a commit error.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

<i>Example: Configuring BGP Route Target Filtering for VPNs</i>
<i>Example: Configuring FEC 129 BGP Autodiscovery for VPWS</i>
<i>Configuring EVPN Routing Instances</i>
<i>Configuring Routing Instances on PE Routers in VPNs</i>
Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) 1460
Configuring an MPLS-Based Layer 3 VPN (CLI Procedure) 1512
<i>path-selection</i>

te-link

Syntax

```
te-link te-link-name {
  disable;
  ethernet-vlan;
  interface interface-name {
    disable;
    local-address ip-address;
    remote-address ip-address;
    remote-id id-number;
  }
  local-address ip-address;
  remote-address ip-address;
  remote-id id-number;
}
```

Hierarchy Level

```
[edit protocols link-management],
[edit protocols link-management peer peer-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

ethernet-vlan option introduced in Junos OS Release 14.2.

Description

Represent a collection of physical ports or time slots. Assign a traffic engineering link to the specified network peer.

Options

te-link-name—Name of the collection of physical ports or the name of the time slots.

disable—Disable the traffic engineering link or an interface to a traffic engineering link.

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 LMP Traffic Engineering Links](#) | 1361

traceoptions (Protocols Link Management)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols link-management],
[edit protocols link-management]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support for **hello-packets**, **packets**, and **state** flags added in Junos OS Release 8.1.

Description

Trace options for the LMP protocol.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`.

files number—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000

Default: 2 files

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **all**—Trace all available operations
- **hello-packets**—Trace hello packets on any LMP control channel
- **init**—Output from the initialization messages

- **packets**—Trace all packets other than hello packets on any LMP control channel
- **parse**—Operation of the parser
- **process**—Operation of the general configuration
- **route-socket**—Operation of route socket events
- **routing**—Operation of the routing protocols
- **server**—Server processing operations
- **show**—**show** command servicing operations
- **state**—Trace state transitions of the LMP control channels and traffic engineering links

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Prevent all users from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable log file access for all users.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Tracing LMP Traffic | 1371](#)

Network Management and Monitoring Guide

transit-delay (OSPF)

Syntax

```
transit-delay seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols ospf area area-id peer-interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id interface interface-name],
[edit logical-systems logical-system-name protocols (ospf | ospf3) area area-id virtual-link],
[edit logical-systems logical-system-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area
  area-id interface interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf area area-id interface
  interface-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf area area-id
  virtual-link],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast
  | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit protocols ospf area area-id peer-interface interface-name],
[edit protocols (ospf | ospf3) area area-id interface interface-name],
[edit protocols (ospf | ospf3) area area-id virtual-link],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols ospf area area-id interface interface-name],
[edit routing-instances routing-instance-name protocols ospf area area-id virtual-link],
[edit routing-instances routing-instance-name protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)
  area area-id interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Support for the **realm** statement introduced in Junos OS Release 9.2.

Support for the **realm** statement introduced in Junos OS Release 9.2 for EX Series switches.

Description

Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays.

You should never have to modify the transit delay time.

Options

seconds—Estimated time, in seconds.

Range: 1 through 65,535 seconds

Default: 1 second

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring OSPF Timers

[Configuring RSVP and OSPF for LMP Peer Interfaces](#) | 1369

upstream-label

Syntax

```
upstream-label {  
  vlan-id vlan-id;  
}
```

Hierarchy Level

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
```

Release Information

Statement introduced in Junos OS Release 14.2.

Description

Specify the upstream label for the bidirectional label-switched path (LSP).

Options

vlan-id *vlan-id*—VLAN ID to be used for the Generalized MPLS (GMPLS) VLAN LSP at the ingress provider edge (PE) to customer edge (CE) link.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring MPLS LSPs for GMPLS](#) | 1372

vrf-target

Syntax

```
vrf-target {
  community;
  auto
  import community-name;
  export community-name;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols l2vpn mesh-group
  mesh-group-name],
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols vpls mesh-group
  mesh-group-name],
[edit routing-instances routing-instance-name protocols evpn vni-options],
[edit routing-instances routing-instance-name protocols l2vpn mesh-group mesh-group-name],
[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name],
[edit switch-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 11.1 for EX Series switches.

Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches. **auto option** was also added at this time.

auto option added in Junos OS Release 19.1R1 for MX series.

Statement introduced in cRPD Release 19.4R1.

Description

Specify a virtual routing and forwarding (VRF) target community. If you configure the **community** option only, default VRF import and export policies are generated that accept and tag routes with the specified target community. The purpose of the **vrf-target** statement is to simplify the configuration by allowing you to configure most statements at the **[edit routing-instances]** hierarchy level. In effect, this statement configures a single policy for import and a single policy for export to replace the per-VRF policies for every community.

You can still create more complex policies by explicitly configuring VRF import and export policies using the **import** and **export** options.

Options

community—Community name.

auto—Automatically derives the route target (RT). The auto-derived route targets have higher precedence over manually configured RT in vrf-target, vrf-export policies, and vrf-import policies.

NOTE: Auto-derived route targets are supported only in virtual switch and EVPN routing instances.

import community-name—Allowed communities accepted from neighbors.

export community-name—Allowed communities sent to neighbors.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Configuring Policies for the VRF Table on PE Routers in VPNs

Example: Configuring FEC 129 BGP Autodiscovery for VPWS

PCEP Configuration Statements

IN THIS CHAPTER

- [pcep | 3072](#)
- [delegation-cleanup-timeout | 3073](#)
- [delegation-priority | 3074](#)
- [destination-ipv4-address | 3075](#)
- [destination-port | 3076](#)
- [label-switched-path-template | 3077](#)
- [lsp-cleanup-timer | 3078](#)
- [lsp-external-controller | 3080](#)
- [max-unknown-messages | 3081](#)
- [max-unknown-requests | 3082](#)
- [message-rate-limit | 3083](#)
- [pce | 3084](#)
- [pce-group \(PCE\) | 3088](#)
- [pce-group \(Protocols PCEP\) | 3089](#)
- [pce-type | 3090](#)
- [querier \(performance-monitoring\) | 3091](#)
- [traceoptions \(PCE\) | 3093](#)
- [traceoptions \(Protocols PCEP\) | 3095](#)
- [update-rate-limit | 3097](#)

pcep

Syntax

```
pcep {
  message-rate-limit messages-per-minute;
  pce pce-id;
  pce-group pce-group-id;
  traceoptions;
  update-rate-limit updates-per-minute;
}
```

Hierarchy Level

[edit protocols]

Release Information

Statement introduced in Junos OS Release 12.3.

Statement introduced in Junos OS Release 16.3R for QFX Series switches.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Configure the Path Computation Client (PCC) parameters.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697](#)

[Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE | 1717](#)

[Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support of PCE-Initiated Point-to-Point LSPs | 1734](#)

[Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Controlled Point-to-Multipoint LSPs | 1750](#)

delegation-cleanup-timeout

Syntax

```
delegation-cleanup-timeout seconds;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Specify the amount of time (in seconds) that a Path Computation Client (PCC) must wait before returning control of all LSPs to the routing protocol process after a PCEP session with the main active stateful Path Computation Element (PCE) is disconnected.

NOTE: In compliance with *draft-ietf-pce-stateful-pce-09*, revoking of PCE-initiated LSP delegations by a PCC happens in a make-before-break fashion before the LSPs are redelegated to an alternate PCE. Starting in Junos OS Release 18.1R1, the **lsp-cleanup-timer** must be greater than or equal to the **delegation-cleanup-timeout** for the PCC to revoke the LSP delegations. If not, the redelegation timeout interval for the PCC can be set to infinity, where the LSP delegations to that PCE remain intact until specific action is taken by the PCC to change the parameters set by the PCE.

Options

seconds—Time (in seconds) that a PCC must wait before returning control of all LSPs to the routing protocol process after a PCEP session with the main active stateful PCE is disconnected.

A value of 0 indicates immediate delegation cleanup.

Range: 0 through 2147483647 seconds

Default: 30 seconds

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pce](#) | [3084](#)

delegation-priority

Syntax

```
delegation-priority priority-number;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Specify the priority number of the active stateful Path Computation Element (PCE). This value is used by the Path Computation Client (PCC) to elect a PCE to delegate LSPs. No two PCEs can have the same delegation-priority value. The PCC elects the PCE with a lower priority as the main active stateful PCE to delegate LSPs.

Options

priority-number—Priority number of the active stateful PCE.

Range: 1 through 65535

Default: 0 (no priority is set)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pce](#) | [3084](#)

destination-ipv4-address

Syntax

```
destination-ipv4-address ipv4-address;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Specify the IPv4 address of the Path Computation Element (PCE) to which the Path Computation Client (PCC) should connect.

Options

ipv4-address—IPv4 address of the PCE.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pce](#) | [3084](#)

destination-port

Syntax

```
destination-port port-number;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFXswitches added in Junos OS Release 16.1R3.

Description

Specify the TCP port number of the Path Computation Element (PCE) to which the Path Computation Client (PCC) should connect.

Options

port-number—Destination TCP port number.

Range: 1 through 65535

Default: 4189

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pce](#) | [3084](#)

label-switched-path-template

Syntax

```
label-switched-path-template {  
    (default-template | lsp-template-name);  
}
```

Hierarchy Level

```
[edit protocols mpls lsp-external-controller lsp-external-controller]
```

Release Information

Statement introduced in Junos OS Release 13.3.

Description

Specify the LSP template with parameters for setting up the PCE-initiated LSPs when the PCE initiating the LSP does not provide the PCE-initiated parameters. When **label-switched-path-template** is not configured, the default LSP parameters are used.

Options

default-template—Specify that the default LSP template be used for the dynamically generated PCE-initiated LSPs.

lsp-template-name—Specify the name of the LSP template to be used for setting up PCE-initiated LSPs.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pcep](#) | [3072](#)

lsp-cleanup-timer

Syntax

```
lsp-cleanup-timer seconds;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]  
[edit protocols pcep pce-group group-id]
```

Release Information

Statement introduced in Junos OS Release 13.3.

Description

Specify the amount of time (in seconds) that the Path Computation Client (PCC) must wait before deleting any non-delegated Path Computation Element (PCE)-initiated LSPs from the failed PCE after a PCEP session terminates.

NOTE: In compliance with *draft-ietf-pce-stateful-pce-09*, revoking of PCE-initiated LSP delegations by a PCC happens in a make-before-break fashion before the LSPs are redelegated to an alternate PCE. Starting in Junos OS Release 18.1R1, the **lsp-cleanup-timer** must be greater than or equal to the **delegation-cleanup-timeout** for the PCC to revoke the LSP delegations. If not, the redelegation timeout interval for the PCC can be set to infinity, where the LSP delegations to that PCE remain intact until specific action is taken by the PCC to change the parameters set by the PCE.

Options

seconds—Time (in seconds) that the PCC must wait before deleting any non-delegated PCE-initiated LSPs from the failed PCE after a PCEP session terminates. Non-delegated PCE-initiated LSPs are deleted immediately.

Range: 0 through 2147483647 seconds

Default: 60

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pcep](#) | [3072](#)

lsp-external-controller

Syntax

```
lsp-external-controller controller-name;
```

Hierarchy Level

```
[edit protocols mpls],  
[edit protocols mpls label-switched-path lsp-name]  
[edit protocols spring-traffic-engineering]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Statement introduced in Junos OS Release 16.1R3 for QFX Series switches.

Support for ACX Series added in Junos OS Release 17.1.

Support at the **[edit protocols spring-traffic-engineering]** hierarchy level introduced in Junos OS Release 17.2.

Description

Enable external path computing capability for the device.

Options

controller-name—Name of the external path computing entity. By default, *pccd* is the only allowed LSP external controller.

Values: *pccd*

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pcep | 3072](#)

[Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697](#)

[Support of SPRING-TE for the Path Computation Element Protocol Overview | 1777](#)

[Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs | 1782](#)

max-unknown-messages

Syntax

```
max-unknown-messages messages-per-minute;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Specify the number of unknown messages per minute that the Path Computation Client (PCC) can receive at maximum after which the PCEP session is closed.

Options

messages-per-minute—Number of unknown messages per minute that the PCC can receive at maximum after which the PCEP session is closed. Recommended value is 5. If the number of unknown messages received by the PCC is greater than or equal to the maximum number, the PCEP session is closed.

Range: 1 through 16384

Default: 0 (disabled or no limit)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pce](#) | [3084](#)

max-unknown-requests

Syntax

```
max-unknown-requests requests-per-minute;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]  
[edit protocols pcep pce-group group-id]
```

Release Information

Statement introduced in Junos OS Release 13.3.

Description

Specifies the number of unknown requests per minute that the Path Computation Client (PCC) can receive at maximum after which the PCEP session is terminated.

Options

requests-per-minute—Number of unknown requests per minute that the PCC can receive at maximum after which the PCEP session is terminated.

Range: 0 through 16384 (0 disables this statement)

Default: 5

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pce](#) | [3084](#)

message-rate-limit

Syntax

```
message-rate-limit messages-per-minute;
```

Hierarchy Level

```
[edit protocols pcep]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Specify the number of messages per minute that the Path Computation Client (PCC) can receive at maximum.

Options

messages-per-minute—Number of messages per minute that the PCC can receive at maximum.

Range: 1 through 16384

Default: 0 (disabled or no limit)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pcep](#) | [3072](#)

pce

Syntax

```
pce pce-id {
    authentication-key key;
    authentication-key-chain key-chain;
    delegation-cleanup-timeout seconds;
    delegation-priority priority-number;
    destination-ipv4-address ipv4-address;
    destination-port port-number;
    local-address ip-address;
    lsp-cleanup-timer seconds;
    lsp-provisioning;
    lsp-retry-delegation;
    lsp-retry-delegation-timer seconds;
    max-sid-depth max-sid-depth;
    max-unknown-messages messages-per-minute;
    max-unknown-requests requests-per-minute;
    p2mp-lsp-init-capability;
    p2mp-lsp-report-capability;
    p2mp-lsp-update-capability;
    pce-group pce-group-name;
    pce_traffic_steering;
    pce-type ;
    request-timer seconds;
    request-priority priority;
    spring-capability;
    traceoptions ;
}
```

Hierarchy Level

```
[edit protocols pcep]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

lsp-cleanup-timer, **lsp-provisioning**, **max-unknown-requests**, **request-timer**, and **request-priority** options introduced in Junos OS Release 13.3.

authentication-key key, **authentication-key-chain key-chain**, and **p2mp-lsp-report-capability** options introduced in Junos OS Release 16.1.

max-sid-depth and **spring-capability** options introduced in Junos OS Release 17.2.

p2mp-lsp-init-capability and **p2mp-lsp-update-capability** options introduced in Junos OS Release 18.3R1 on all platforms.

pce_traffic_steering option introduced in Junos OS Release 19.4R1 on all platforms.

Description

Configure per-Path Computation Element (PCE) parameters.

Options

pce-id—IP address of the PCE.

authentication-key key—(Optional) Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

It is recommended to define and bind an authentication key for securing a PCEP session, as opposed to binding an authentication keychain.

authentication-key-chain key-chain—(Optional) Authentication keychain password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

local-address ip-address—(Optional) IP address of the local end of the PCEP session, the PCC.

lsp-retry-delegation—(Optional) Enable retry LSP delegation process.

lsp-retry-delegation-timer—(Optional) Specify the amount of time (in seconds) that the Path Computation Client (PCC) must wait before retrying delegation of Path Computation Element (PCE)-initiated LSPs in case of delegation failure or re-delegation.

Default: 3600 seconds

Range: 0 through 4294967294 seconds

max-sid-depth—(Optional) Specify the maximum value for service identifier (SID) depth.

Default: 5

Range: 1 through 5

p2mp-lsp-init-capability—(Optional) Capability to provision point-to-multipoint RSVP-TE LSPs by a PCE. By default, this capability is not supported on a PCC, and should be explicitly configured to enable PCE-initiated point-to-multipoint LSPs.

p2mp-lsp-report-capability—(Optional) Capability to report point-to-multipoint RSVP-TE LSPs to a PCE. By default, this capability is not supported on a PCC, and should be explicitly configured to enable reporting of point-to-multipoint LSPs to a PCE.

p2mp-lsp-update-capability—(Optional) Capability to update point-to-multipoint RSVP-TE LSP parameters by a PCE. By default, this capability is not supported on a PCC, and should be explicitly configured to enable updating of PCE-initiated point-to-multipoint LSPs.

pce_traffic_steering—(Optional) Configure the flow specification capability (also called traffic steering functionality) for enabling the mapping of PCE-initiated point-to-multipoint LSPs to an MVPN routing-instance.

spring-capability—(Optional) Enable SPRING-based provisioning for the PCE.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pcep | 3072](#)

[Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs | 1782](#)

[Support of the Path Computation Element Protocol for RSVP-TE Overview | 1697](#)

pce-group (PCE)

Syntax

```
pce-group pce-group-name;
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Specify the Path Computation Element (PCE) group to which the configured PCE belongs.

Options

pce-group-name—Name of the PCE group.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pce](#) | [3084](#)

pce-group (Protocols PCEP)

Syntax

```
pce-group pce-group-id {
  delegation-cleanup-timeout seconds;
  max-unknown-messages messages-per-minute;
  pce-type {
    active stateful;
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag (all | pcep);
    no-remote-trace;
  }
}
```

Hierarchy Level

[edit protocols pcep]

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Configure the Path Computation Element (PCE) group parameters. A maximum of 10 PCE groups can be configured at any given point in time.

The remaining statements are explained separately.

NOTE: A PCE group can include PCEs that are either only stateful or only active stateful. A combination of stateful PCEs and active stateful PCEs in one PCE group is not supported.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pcep](#) | [3072](#)

pce-type

Syntax

```
pce-type {  
    active stateful;  
}
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Configure the path computation element (PCE) type:

- **active**—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegated control over their LSPs to the PCE.
- **stateful**—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update the LSP state. A PCC maintains synchronization with the PCE.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[pce](#) | [3084](#)

querier (performance-monitoring)

Syntax

```

querier {
  delay {
    traffic-class tc-value {
      average-sample-size sample size;
      padding-size size;
      query-interval milliseconds;
      rtt-delay-threshold rtt threshold value;
      twcd-delay-threshold twcd threshold value;
    }
  }
  loss {
    traffic-class tc-value {
      average-sample-size sample size;
      loss-threshold loss threshold value;
      loss-threshold-window number of samples for loss threshold;
      measurement-quantity bytes|packets;
      query-interval milliseconds;
    }
  }
  loss-delay {
    traffic-class tc-value {
      average-sample-size sample size;
      loss-threshold loss threshold value;
      loss-threshold-window number of samples for loss threshold;
      measurement-quantity bytes|packets;
      padding-size size;
      query-interval milliseconds;
      rtt-delay-threshold rtt threshold value;
      twcd-delay-threshold twcd threshold value;
    }
  }
}

```

Hierarchy Level

```

[edit protocols mpls oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring],
[edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring]

```


Release Information

Statement introduced in Junos OS Release 15.1.

Command introduced in Junos OS Release 16.1R3 for QFX Series switches.

Description

Configure querier options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

traceoptions (PCE)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag (all | pcep);
  no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols pcep pce pce-id]
```

Description

Configure the Path Computation Element Protocol (PCEP) tracing options.

Options

filename—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`.

files number—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000 files

Default: 2 files. If you specify a maximum number of files, you must also include the ***size*** statement to specify the maximum file size.

flag—Area of path computation client process (pccd) to enable debugging output.

- **all**—Trace all areas of PCD code.
- **pcep**—Trace Path Computation Element Protocol (PCEP) operations.

no-remote-trace—(Optional) Disable remote tracing options.

no-world-readable—(Optional) Allow only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named `trace-file` reaches this size, it is renamed `trace-file.0`. When the trace-file again reaches this size, `trace-file.0` is renamed `trace-file.1` and `trace-file` is renamed `trace-file.0`. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: ***xk*** to specify KB, ***xm*** to specify MB, or ***xg*** to specify GB.

Range: 10 KB through the maximum file size supported on your system.

Default: 1 MB. If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pce](#) | [3084](#)

traceoptions (Protocols PCEP)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

[edit protocols pcep]

Description

Configure the Path Computation Element Protocol (PCEP) tracing options.

Options

filename—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`.

files number—(Optional) Maximum number of trace files. When a trace file named `trace-file` reaches its maximum size, it is renamed `trace-file.0`, then `trace-file.1`, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000 files

Default: 2 files. If you specify a maximum number of files, you must also include the ***size*** statement to specify the maximum file size.

flag—Area of path computation client process (pccd) to enable debugging output.

PCEP Tracing Flags

- ***all***—Trace all areas of PCCD code
- ***pccd-config***—All configuration parsing operations
- ***pccd-core***—PCCD core operations
- ***pccd-functions***—PCCD function entries and outs
- ***pccd-main***—PCCD main module
- ***pccd-rpd***—PCCD communication with RPD
- ***pccd-ui***—PCCD user interface handling

no-remote-trace—(Optional) Disable remote tracing options.

no-world-readable—(Optional) Allow only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches this size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.

Range: 10 KB through the maximum file size supported on your system.

Default: 1 MB. If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.

routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pcep](#) | [3072](#)

update-rate-limit

Syntax

```
update-rate-limit updates-per-minute;
```

Hierarchy Level

```
[edit protocols pcep]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for PTX Series added in Junos OS Release 14.2.

Support for QFX Series switches added in Junos OS Release 16.1R3.

Support for ACX Series added in Junos OS Release 17.1R1.

Description

Specify the number of updates per minute that the Path Computation Client (PCC) can receive at maximum. Updates above this limit are ignored by the PCC.

Options

updates-per-minute—Number of updates per minute that the PCC can receive at maximum.

Range: 1 through 16384

Default: 0 (disabled or no limit)

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [pcep](#) | [3072](#)

11

PART

Operational Commands

MPLS Operational Commands | **3101**

RSVP Operational Commands | **3355**

LDP Operational Commands | **3429**

CCC and TCC Operational Commands | **3515**

PCEP Operational Commands | **3547**

MPLS Operational Commands

IN THIS CHAPTER

- `clear mpls lsp` | 3103
- `clear mpls container-lsp` | 3105
- `clear performance-monitoring mpls lsp` | 3107
- `monitor mpls delay rsvp` | 3108
- `monitor mpls loss rsvp` | 3114
- `monitor mpls loss-delay rsvp` | 3120
- `ping mpls bgp` | 3125
- `ping mpls lsp-end-point` | 3128
- `ping mpls l2circuit` | 3131
- `ping mpls l2vpn` | 3134
- `ping mpls l3vpn` | 3137
- `request mpls container-lsp` | 3140
- `request mpls lsp adjust-autobandwidth` | 3142
- `show connections` | 3144
- `show dynamic-tunnels database` | 3148
- `show link-management` | 3153
- `show link-management peer` | 3157
- `show link-management routing` | 3160
- `show link-management statistics` | 3164
- `show link-management te-link` | 3167
- `show mpls abstract-hop-membership` | 3170
- `show mpls admin-groups` | 3172
- `show mpls association` | 3174
- `show mpls call-admission-control` | 3176
- `show mpls container-lsp` | 3179
- `show mpls context-identifier` | 3188
- `show mpls correlation label` | 3191
- `show mpls correlation nexthop-id` | 3192

- [show mpls cspf | 3194](#)
- [show mpls diffserv-te | 3197](#)
- [show mpls interface | 3199](#)
- [show mpls egress-protection | 3201](#)
- [show mpls interface | 3204](#)
- [show mpls label usage | 3207](#)
- [show mpls label usage label-range | 3211](#)
- [show mpls lsp | 3214](#)
- [show mpls lsp abstract-computation | 3240](#)
- [show mpls lsp autobandwidth | 3243](#)
- [show mpls path | 3246](#)
- [show mpls srlg | 3248](#)
- [show mpls static-lsp | 3250](#)
- [show performance-monitoring mpls lsp | 3254](#)
- [show route forwarding-table | 3262](#)
- [show route table | 3273](#)
- [show ted database | 3329](#)
- [show ted link | 3341](#)
- [show ted protocol | 3346](#)
- [traceroute mpls bgp | 3348](#)
- [transit \(Chained Composite Next Hops\) | 3352](#)

clear mpls lsp

List of Syntax

[Syntax on page 3103](#)

[Syntax \(EX and QFX Series Switches\) on page 3103](#)

Syntax

```
clear mpls lsp
<all>
<autobandwidth>
<counters>
<logical-system (all | logical-system-name)>
<name name>
<optimize | optimize-aggressive>
<path regular-expression>
<statistics>
```

Syntax (EX and QFX Series Switches)

```
clear mpls lsp
<all>
<autobandwidth>
<name name>
<optimize | optimize-aggressive>
<path regular-expression>
<statistics>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.



CAUTION: This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.

Options

all—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.

autobandwidth—(Optional) Clear LSP autobandwidth counters.

counters—(Optional) Reset the flap and the MBB counters to zero.

logical-system (**all** | **logical-system-name**)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name **name**—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos Network Interfaces Configuration Guide*.

optimize | **optimize-aggressive**—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.

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 | 3214](#)

[show rsvp session | 3390](#)

List of Sample Output

[clear mpls lsp all on page 3104](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear mpls lsp all
```

```
user@host> clear mpls lsp all
```


clear mpls container-lsp

Syntax

```
clear mpls container-lsp
<autobandwidth>
<history>
<logical-system (all | logical-system-name)>
<member>
<name name>
<optimize | optimize-aggressive>
<statistics>
```

Release Information

Statement introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

Description

Release the routes and states associated with MPLS container label-switched paths (LSPs), and start new LSPs.

Options

none—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

autobandwidth—(Optional) Clear LSP autobandwidth counters.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name *name*—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos Network Interfaces Configuration Guide*.

optimize | optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.

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 container-lsp | 3179](#)

[request mpls container-lsp | 3140](#)

List of Sample Output

[clear mpls container-lsp on page 3106](#)

[clear mpls container-lsp name on page 3106](#)

[clear mpls container-lsp statistics on page 3106](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mpls container-lsp

```
user@host> clear mpls container-lsp
```

clear mpls container-lsp name

```
user@host> clear mpls container-lsp name name
```

clear mpls container-lsp statistics

```
user@host> clear mpls container-lsp statistics
```


clear performance-monitoring mpls lsp

Syntax

```
clear performance-monitoring mpls lsp  
<name lsp-name>
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Restart the performance monitoring statistics.

Options

none—Reset and restart all performance monitoring for all LSPs.

name *lsp-name*—(Optional) Reset and restart performance monitoring for the specified LSP.

Required Privilege Level

clear

RELATED DOCUMENTATION

[performance-monitoring \(Protocols MPLS\) | 2727](#)

[show performance-monitoring mpls lsp | 3254](#)

List of Sample Output

[clear performance-monitoring mpls lsp on page 3107](#)

Output Fields

When you enter this command, performance monitoring is restarted.

Sample Output

```
clear performance-monitoring mpls lsp
```

```
user@host> clear performance-monitoring mpls lsp
```


monitor mpls delay rsvp

Syntax

```
monitor mpls delay rsvp lsp-name
<detail>
<count count>
<interval seconds>
<padding-size padding-size>
<traffic-class traffic-class>
```

Release Information

Command introduced in Junos OS Release 14.2.

Description

Perform an on-demand delay measurement and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options

lsp-name—Name of the associated bidirectional MPLS UHP LSP for which the delay measurement is performed.

detail—(Optional) Display detailed output of the LSP delay measurement.

count count—(Optional) Specify the number of delay measurements to be carried out for the MPLS UHP LSP. For LSP delay measurement, the number of queries sent is the specified count number plus one additional query, because the LSP delay is measured using successive messages.

Default: 10

Range: 1 through 1000000

interval seconds—(Optional) Specify in seconds the interval between two successive query messages.

Range: 1 through 255 seconds

padding-size padding-size— (Optional) Specify the length of padding TLV to be included in the query message.

Range: 0 through 1500

traffic-class traffic-class—(Optional) Specify the traffic class for the LSP delay measurement. When the traffic-class value is not specified, the default traffic-class code-point of 111 is used.

Range: 0 though 7

Required Privilege Level

view

RELATED DOCUMENTATION

[monitor mpls loss rsvp | 3114](#)

[monitor mpls loss-delay rsvp | 3120](#)

[Example: Configuring On-Demand Loss and Delay Measurement | 250](#)

List of Sample Output
[monitor mpls lsp delay rsvp count on page 3110](#)
[monitor mpls lsp delay rsvp count detail on page 3111](#)
Output Fields

[Table 81 on page 3109](#) describes the output fields for the **monitor mpls delay rsvp** command. Output fields are listed in the approximate order in which they appear.

Table 81: monitor mpls delay rsvp Output Fields

Field Name	Field Description	Level of Output
Current two-way channel delay	Sum of packet delays, excluding the processing time of the remote provider edge (PE) router.	All Levels
Current round-trip-time	Total time taken for completing round-trip of packet.	All Levels
Best two-way channel delay	Best available two-way channel delay count.	All Levels
Worst two-way channel delay	Worst available two-way channel delay count.	All Levels
Average two-way channel delay	Average of the available two-way channel delay counts.	All Levels
Best round-trip-time	Best available round-trip-time count.	All Levels
Worst round-trip-time	Worst available round-trip-time count.	All Levels
Average round-trip-time	Average of the available round-trip-time counts.	All Levels
Average forward delay variation	Average of the variation in forward delay.	All Levels
Average reverse delay variation	Average of the variation in reverse delay.	All Levels

Table 81: monitor mpls delay rsvp Output Fields (continued)

Field Name	Field Description	Level of Output
DM queries sent	Number of queries sent for delay measurement.	All Levels
DM responses received	Number of responses received for delay measurement queries.	All Levels
DM queries timedout	Number of timed out queries sent for delay measurement.	All Levels
DM responses dropped due to errors	Number of loss measurement responses dropped due to errors.	All Levels
Response code	Status of the messages used for delay measurement. Response code can be one of the following: <ul style="list-style-type: none"> • Success—Successful response code. • Failed—Failed response code. 	detail
Querier transmit timestamp	Timestamp on the query message when the message is sent out the ingress PE router (querier). This is done in the hardware before packet is sent out of an interface.	detail
Responder receive timestamp	Timestamp on the response message when the message is received by the egress PE router (responder). This is done in the hardware before packet is received by an interface.	detail
Responder transmit timestamp	Timestamp on the query message when the message is sent out the egress PE router (responder). This is done in the hardware before packet is sent out of an interface.	detail
Querier receive timestamp	Timestamp on the response message when the message is received by the ingress PE router (querier). This is done in the hardware before packet is received by an interface.	detail

Sample Output

```
monitor mpls lsp delay rsvp count
```

```
user@host> monitor mpls lsp delay rsvp LSP-A count 2
```



```

(1)
Current two-way channel delay      : 44 usecs
Current round-trip-time            : 3243 usecs
(2)
Current two-way channel delay      : 45 usecs
Current round-trip-time            : 1752 usecs

Best two-way channel delay         : 44 usecs
Worst two-way channel delay        : 45 usecs
Average two-way channel delay      : 45 usecs
Best round-trip-time               : 1752 usecs
Worst round-trip-time              : 3243 usecs
Average round-trip-time            : 2498 usecs
Average forward delay variation    : 1 usecs
Average reverse delay variation    : 1 usecs

DM queries sent                    : 2
DM responses received              : 2
DM queries timedout                : 0
DM responses dropped due to errors : 0

```

monitor mpls lsp delay rsvp count detail

user@host> monitor mpls lsp delay rsvp LSP-A count 2 detail

```

(1)
Response code                      : Success
Querier transmit timestamp         : 1404129122 secs, 479955401 nsecs
Responder receive timestamp        : 1404129122 secs, 468519022 nsecs
Responder transmit timestamp       : 1404129122 secs, 470255123 nsecs
Querier receive timestamp          : 1404129122 secs, 481736403 nsecs
Current two-way channel delay      : 44 usecs
Current round-trip-time            : 1781 usecs
(2)
Response code                      : Success
Querier transmit timestamp         : 1404129123 secs, 480926210 nsecs
Responder receive timestamp        : 1404129123 secs, 469488696 nsecs
Responder transmit timestamp       : 1404129123 secs, 471130706 nsecs
Querier receive timestamp          : 1404129123 secs, 482613911 nsecs
Current two-way channel delay      : 45 usecs
Current round-trip-time            : 1687 usecs

```



```

Best two-way channel delay      : 44 usecs
Worst two-way channel delay     : 45 usecs
Average two-way channel delay   : 45 usecs
Best round-trip-time           : 1687 usecs
Worst round-trip-time          : 1781 usecs
Average round-trip-time        : 1734 usecs
Average forward delay variation : 1 usecs
Average reverse delay variation : 1 usecs

DM queries sent                 : 2
DM responses received           : 2
DM queries timedout             : 0
DM responses dropped due to errors : 0
user@host> monitor mpls loss-delay-measurement lsp LSP1_A_to_B count 2
(1)
Current forward loss            : 0 packets
Current forward loss ratio      : 0.000000
Current forward throughput      : 0.957 kpps
Current reverse loss            : 0 packets
Current reverse loss ratio      : 0.000000
Current reverse throughput      : 0.962 kpps
Current two-way channel delay    : 48 usecs
Current round-trip-time         : 3476 usecs
(2)
Current forward loss            : 0 packets
Current forward loss ratio      : 0.000000
Current forward throughput      : 0.599 kpps
Current reverse loss            : 0 packets
Current reverse loss ratio      : 0.000000
Current reverse throughput      : 0.599 kpps
Current two-way channel delay    : 50 usecs
Current round-trip-time         : 1856 usecs

Cumulative forward transmit count : 1557
Cumulative forward loss            : 0 packets
Average forward loss ratio         : 0.000000
Average forward throughput         : 0.778 kpps
Cumulative reverse transmit count  : 1562
Cumulative reverse loss            : 0 packets
Average reverse loss ratio         : 0.000000
Average reverse throughput         : 0.780 kpps

Best two-way channel delay      : 48 usecs
Worst two-way channel delay     : 50 usecs

```


Average two-way channel delay	: 49 usecs
Best round-trip-time	: 1856 usecs
Worst round-trip-time	: 3476 usecs
Average round-trip-time	: 2445 usecs
Average forward delay variation	: 1 usecs
Average reverse delay variation	: 1 usecs
LDM queries sent	: 3
LDM responses received	: 3
LDM queries timedout	: 0
LDM responses dropped due to errors	: 0

monitor mpls loss rsvp

Syntax

```
monitor mpls loss rsvp lsp-name
<detail>
<bytes>
<count count>
<interval seconds>
<traffic-class traffic-class>
```

Release Information

Command introduced in Junos OS Release 14.2.

Description

Perform an on-demand loss measurement and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options

lsp-name—Name of the associated bidirectional MPLS UHP LSP for which the loss measurement is performed.

detail—(Optional) Display detailed output of the LSP loss measurement.

bytes—(Optional) Specify the measurement quantity for the LSP loss measurement as bytes. By default, LSP loss is measured in packets.

NOTE: The byte count of a packet sent or received over a channel counts only the payload, including the total length of that packet and excluding the headers, labels, and framing of the channel itself.

count count—(Optional) Specify the number of loss measurements to be carried out for the MPLS UHP LSP. For LSP loss measurement, the number of queries sent is the specified count number plus one additional query, because the LSP loss is measured using successive messages.

Default: 10

Range: 1 through 1000000

interval seconds—(Optional) Specify in seconds the interval between two successive query messages.

Range: 1 through 255 seconds

traffic-class *traffic-class*—(Optional) Specify the traffic class and enable traffic-class-statistics for the LSP loss measurement.

Range: 0 though 7

Required Privilege Level

view

RELATED DOCUMENTATION

[monitor mpls delay rsvp | 3108](#)

[monitor mpls loss-delay rsvp | 3120](#)

[Example: Configuring On-Demand Loss and Delay Measurement | 250](#)

List of Sample Output

[monitor mpls lsp loss rsvp count on page 3117](#)

[monitor mpls lsp loss rsvp detail on page 3118](#)

Output Fields

[Table 81 on page 3109](#) describes the output fields for the **monitor mpls loss rsvp** command. Output fields are listed in the approximate order in which they appear.

Table 82: monitor mpls loss rsvp Output Fields

Field Name	Field Description	Level of Output
Current forward loss	Difference between the current forward transmit count and the current forward receive count.	All Levels
Current forward loss ratio	Total packet loss (current forward loss divided by current forward transmit count).	All Levels
Current forward throughput	Current forward transmit count divided by 1000.	All Levels
Current reverse loss	Difference between the current reverse transmit count and the current reverse receive count.	All Levels
Current reverse loss ratio	Total packet loss (current reverse loss divided by current reverse transmit count).	All Levels
Current reverse throughput	Current reverse transmit count divided by 1000.	All Levels

Table 82: monitor mpls loss rsvp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Cumulative forward transmit count	Cumulative forward transmit counter value at the time the loss measurement message was originated.	All Levels
Cumulative forward loss	Cumulative forward loss counter value at the time the loss measurement message was originated.	All Levels
Average forward loss ratio	Average packet loss (current forward loss divided by current forward transmit count).	All Levels
Average forward throughput	Average forward transmit count divided by 1000.	All Levels
Cumulative reverse transmit count	Cumulative reverse transmit counter value at the time the loss measurement message was originated.	All Levels
Cumulative reverse loss	Difference between the cumulative reverse transmit count and the cumulative reverse receive count.	All Levels
Average reverse loss ratio	Average packet loss (average reverse loss divided by average reverse transmit count).	All Levels
Average reverse throughput	Average reverse transmit count divided by 1000.	All Levels
LM queries sent	Number of queries sent for loss measurement.	All Levels
LM responses received	Number of responses received for loss measurement queries.	All Levels
LM queries timedout	Number of timed out queries sent for loss measurement.	All Levels
LM responses dropped due to errors	Number of loss measurement responses dropped due to errors.	All Levels
Response code	Status of the messages used for loss measurement. Response code can be one of the following: <ul style="list-style-type: none"> ● Success—Successful response code. ● Failed—Failed response code. 	detail
Origin timestamp	Time and date the loss measurement message is originated without any specific format (NTP and PTP).	detail

Table 82: monitor mpls loss rsvp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Forward transmit count	Forward transmit counter value at the time the loss measurement message was originated.	detail
Forward receive count	Forward receive counter value at the time the loss measurement message was originated.	detail
Reverse transmit count	Reverse transmit counter value at the time the loss measurement message was originated.	detail
Reverse receive count	Reverse receive counter value at the time the loss measurement message was originated.	detail
Current forward transmit count	Difference between the current forward transit count and the previous forward transit count.	detail
Current forward receive count	Difference between the current forward receive count and the previous forward receive count.	detail
Current reverse transmit count	Difference between the current reverse transit count and the previous reverse transit count.	detail
Current reverse receive count	Difference between the current reverse receive count and the previous reverse receive count.	detail

Sample Output

monitor mpls lsp loss rsvp count

user@host> **monitor mpls lsp loss rsvp count 2**

```
(1)
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 1.006 kpps
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 1.007 kpps
```



```

(2)
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 0.559 kpps
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 0.562 kpps

Cumulative forward transmit count : 1559
Cumulative forward loss           : 0 packets
Average forward loss ratio       : 0.000000
Average forward throughput       : 0.782 kpps
Cumulative reverse transmit count : 1563
Cumulative reverse loss           : 0 packets
Average reverse loss ratio       : 0.000000
Average reverse throughput       : 0.784 kpps

LM queries sent                 : 3
LM responses received           : 3
LM queries timedout             : 0
LM responses dropped due to errors : 0

```

monitor mpls lsp loss rsvp detail

user@host> monitor mpls lsp loss rsvp detail

```

(0)
Response code                   : Success
Origin timestamp                 : 1404129082 secs, 905571890 nsecs
Forward transmit count           : 83040
Forward receive count            : 83040
Reverse transmit count           : 83100
Reverse receive count            : 83100

(1)
Response code                   : Success
Origin timestamp                 : 1404129083 secs, 905048410 nsecs
Forward transmit count           : 83841
Forward receive count            : 83841
Reverse transmit count           : 83904
Reverse receive count            : 83904
Current forward transmit count   : 801
Current forward receive count    : 801
Current forward loss             : 0 packets
Current forward loss ratio      : 0.000000

```



```

Current forward throughput      : 0.801 kpps
Current reverse transmit count  : 804
Current reverse receive count   : 804
Current reverse loss            : 0 packets
Current reverse loss ratio      : 0.000000
Current reverse throughput      : 0.804 kpps
(2)
Response code                   : Success
Origin timestamp                : 1404129084 secs, 904828715 nsecs
Forward transmit count          : 84423
Forward receive count           : 84423
Reverse transmit count          : 84487
Reverse receive count           : 84487
Current forward transmit count  : 582
Current forward receive count   : 582
Current forward loss            : 0 packets
Current forward loss ratio      : 0.000000
Current forward throughput      : 0.582 kpps
Current reverse transmit count  : 583
Current reverse receive count   : 583
Current reverse loss            : 0 packets
Current reverse loss ratio      : 0.000000
Current reverse throughput      : 0.583 kpps

Cumulative forward transmit count : 1383
Cumulative forward loss           : 0 packets
Average forward loss ratio        : 0.000000
Average forward throughput        : 0.692 kpps
Cumulative reverse transmit count : 1387
Cumulative reverse loss           : 0 packets
Average reverse loss ratio        : 0.000000
Average reverse throughput        : 0.694 kpps

LM queries sent                  : 3
LM responses received            : 3
LM queries timedout              : 0
LM responses dropped due to errors : 0

```


monitor mpls loss-delay rsvp

Syntax

```
monitor mpls loss-delay rsvp lsp-name
<detail>
<bytes>
<count count>
<interval seconds>
<padding-size padding-size>
<traffic-class traffic-class>
```

Release Information

Command introduced in Junos OS Release 14.2.

Description

Perform a simultaneous on-demand loss and delay measurement using combined loss and delay messages, and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options

lsp-name—Name of the associated bidirectional MPLS UHP LSP for which the delay measurement is performed.

detail—(Optional) Display detailed output of the LSP delay measurement.

bytes—(Optional) Specify the measurement quantity for the LSP loss measurement as bytes. By default, LSP loss is measured in packets.

NOTE: The byte count of a packet sent or received over a channel counts only the payload, including the total length of that packet and excluding the headers, labels, and framing of the channel itself.

count count—(Optional) Specify the number of delay measurements to be carried out for the MPLS UHP LSP. For LSP delay measurement, the number of queries sent is the specified count number plus one additional query, because the LSP delay is measured using successive messages.

Default: 10

Range: 1 through 1000000

interval seconds—(Optional) Specify in seconds the interval between two successive query messages.

Range: 1 through 255 seconds

padding-size *padding-size*—(Optional) Specify the length of padding TLV to be included in the query message.

Range: 0 through 1500

traffic-class *traffic-class*—(Optional) Specify the traffic class for the LSP delay measurement. When the traffic-class value is not specified, the default traffic-class code-point of 111 is used.

Range: 0 though 7

Required Privilege Level

view

RELATED DOCUMENTATION

[monitor mpls loss rsvp | 3114](#)

[monitor mpls delay rsvp | 3108](#)

[Example: Configuring On-Demand Loss and Delay Measurement | 250](#)

List of Sample Output

[monitor mpls loss-delay rsvp count on page 3121](#)

[monitor mpls loss-delay rsvp count detail on page 3122](#)

Output Fields

For output field descriptions, see the [monitor mpls loss rsvp](#) and [monitor mpls delay rsvp](#) commands.

Sample Output

monitor mpls loss-delay rsvp count

user@host> **monitor mpls loss-delay rsvp LSP-A count 2**

```
(1)
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 0.957 kpps
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 0.962 kpps
Current two-way channel delay   : 48 usecs
Current round-trip-time        : 3476 usecs
(2)
```



```

Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 0.599 kpps
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 0.599 kpps
Current two-way channel delay   : 50 usecs
Current round-trip-time        : 1856 usecs

Cumulative forward transmit count : 1557
Cumulative forward loss           : 0 packets
Average forward loss ratio       : 0.000000
Average forward throughput       : 0.778 kpps
Cumulative reverse transmit count : 1562
Cumulative reverse loss          : 0 packets
Average reverse loss ratio       : 0.000000
Average reverse throughput       : 0.780 kpps

Best two-way channel delay       : 48 usecs
Worst two-way channel delay      : 50 usecs
Average two-way channel delay    : 49 usecs
Best round-trip-time             : 1856 usecs
Worst round-trip-time            : 3476 usecs
Average round-trip-time          : 2445 usecs
Average forward delay variation  : 1 usecs
Average reverse delay variation  : 1 usecs

LDM queries sent                 : 3
LDM responses received           : 3
LDM queries timedout             : 0
LDM responses dropped due to errors : 0

```

monitor mpls loss-delay rsvp count detail

user@host> monitor mpls loss-delay rsvp LSP-A count 2 detail

```

(0)
Response code           : Success
Forward transmit count   : 142049
Forward receive count    : 142049
Reverse transmit count   : 142167
Reverse receive count    : 142167
Querier transmit timestamp : 1404129161 secs, 554422723 nsecs

```



```

Responder receive timestamp      : 1404129161 secs, 542877570 nsecs
Responder transmit timestamp     : 1404129161 secs, 546004545 nsecs
Querier receive timestamp       : 1404129161 secs, 557599327 nsecs
(1)
Response code                   : Success
Forward transmit count          : 143049
Forward receive count           : 143049
Reverse transmit count          : 143168
Reverse receive count           : 143168
Current forward transmit count   : 1000
Current forward receive count    : 1000
Current forward loss             : 0 packets
Current forward loss ratio       : 0.000000
Current forward throughput       : 1.000 kpps
Current reverse transmit count   : 1001
Current reverse receive count    : 1001
Current reverse loss             : 0 packets
Current reverse loss ratio       : 0.000000
Current reverse throughput       : 1.001 kpps
Querier transmit timestamp       : 1404129162 secs, 554465742 nsecs
Responder receive timestamp      : 1404129162 secs, 542919166 nsecs
Responder transmit timestamp     : 1404129162 secs, 545812736 nsecs
Querier receive timestamp       : 1404129162 secs, 557409175 nsecs
Current two-way channel delay    : 49 usecs
Current round-trip-time         : 2943 usecs
(2)
Response code                   : Success
Forward transmit count          : 143677
Forward receive count           : 143677
Reverse transmit count          : 143799
Reverse receive count           : 143799
Current forward transmit count   : 628
Current forward receive count    : 628
Current forward loss             : 0 packets
Current forward loss ratio       : 0.000000
Current forward throughput       : 0.627 kpps
Current reverse transmit count   : 631
Current reverse receive count    : 631
Current reverse loss             : 0 packets
Current reverse loss ratio       : 0.000000
Current reverse throughput       : 0.630 kpps
Querier transmit timestamp       : 1404129163 secs, 556698575 nsecs
Responder receive timestamp      : 1404129163 secs, 545150128 nsecs
Responder transmit timestamp     : 1404129163 secs, 546918408 nsecs

```


Querier receive timestamp	: 1404129163 secs, 558515047 nsecs
Current two-way channel delay	: 48 usecs
Current round-trip-time	: 1816 usecs
Cumulative forward transmit count	: 1628
Cumulative forward loss	: 0 packets
Average forward loss ratio	: 0.000000
Average forward throughput	: 0.813 kpps
Cumulative reverse transmit count	: 1632
Cumulative reverse loss	: 0 packets
Average reverse loss ratio	: 0.000000
Average reverse throughput	: 0.815 kpps
Best two-way channel delay	: 48 usecs
Worst two-way channel delay	: 49 usecs
Average two-way channel delay	: 49 usecs
Best round-trip-time	: 1816 usecs
Worst round-trip-time	: 3176 usecs
Average round-trip-time	: 2645 usecs
Average forward delay variation	: 1 usecs
Average reverse delay variation	: 0 usecs
LDM queries sent	: 3
LDM responses received	: 3
LDM queries timedout	: 0
LDM responses dropped due to errors	: 0

ping mpls bgp

Syntax

```
ping mpls bgp fec
<bottom-label-ttl>
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<size bytes>
<source source-address>
<sweep>
```

Release Information

Command introduced in Junos OS Release 11.1.

Description

Check the operability of MPLS BGP-signaled label-switched path (LSP) connections. Press Ctrl+c to interrupt a **ping mpls bgp** command.

NOTE: The **ping mpls bgp fec** command only supports single paths.

Options

bottom-label-ttl—(Optional) Time-to-live (TTL) value for the bottom label in the label stack. The range of values is 1 through 255. The default value is **255**.

count count—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination address—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp forwarding-class—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping a BGP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

instance routing-instance-name—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

size bytes—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

NOTE: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

List of Sample Output

[ping mpls bgp fec count on page 3127](#)

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--- lsping statistics ---10 packets transmitted, 3 packets received, 70%  
packet loss 4 packets received with error status, not counted as received.
```


ping mpls lsp-end-point

Syntax

```
ping mpls lsp-end-point prefix-name
<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 before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

The **size** and **sweep** options were introduced in Junos OS Release 9.6.

The **instance** option was introduced in Junos OS Release 10.0.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Check the operability of MPLS label-switched path (LSP) endpoint 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.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

instance *routing-instance-name*—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

prefix-name—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.

size bytes—(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.

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).

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.

NOTE: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

List of Sample Output

[ping mpls lsp-end-point detail on page 3130](#)

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
```


ping mpls l2circuit

Syntax

```
ping mpls l2circuit (interface interface-name | virtual-circuit virtual-circuit-id neighbor address)
<count count>
<destination address>
<detail>
<exp forwarding-class>
<logical-system (all | logical-system-name)>
reply-mode (application-level-control-channel | ip-udp | no-reply)
<size bytes>
<source source-address>
<sweep>
<v1>
```

Release Information

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 12.3X50 for the QFX Series.

The **size** and **sweep** options were introduced in Junos OS Release 9.6.

The **reply-mode** option and its suboptions are introduced in Junos OS Release 10.4R1.

Description

Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a **ping mpls l2circuit** command.

NOTE: This command is not supported on EX4500 and EX4550 switches.

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.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

interface *interface-name*—Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

reply-mode—(Optional) Reply mode for the ping request. This option has the following suboptions:

application-level-control-channel—Reply using an application level control channel.

ip-udp—Reply using an IPv4 or IPv6 UDP packet.

no-reply—Do not reply to the ping request.

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).

v1—(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 3133](#)

[ping mpls l2circuit virtual-circuit detail on page 3133](#)

[ping mpls l2circuit interface <interface-name> reply-mode on page 3133](#)

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 l2vpn

Syntax

```
ping mpls l2vpn (instance instance-name local-site-id local-site-id-number remote-site-id remote-site-id-number |
  interface interface-name)
<bottom-label-ttl>
<count count>
<destination address>
<detail>
<exp forwarding-class>
<logical-system (all | logical-system-name)>
reply-mode (application-level-control-channel | ip-udp | no-reply)
<size bytes>
<source source-address>
<sweep>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

The **size** and **sweep** options were introduced in Junos OS Release 9.6.

The **reply-mode** option and its suboptions are introduced in Junos OS Release 10.4R1.

Description

Check the operability of MPLS Layer 2 virtual private network (VPN) connections. Type Ctrl+c to interrupt a **ping mpls l2vpn** command.

Options

bottom-label-ttl—(Optional) Display the time-to-live value for the bottom label in the label stack.

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.

instance *instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number*—Ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier, testing the integrity of the Layer 2 VPN circuit (specified by the identifiers) between the ingress and egress provider edge (PE) routers or switches.

interface *interface-name*—Ping an interface configured for the Layer 2 VPN on the egress PE router or switch.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

reply-mode—(Optional) Reply mode for the ping request. This option has the following suboptions:

application-level-control-channel—Reply using an application level control channel.

ip-udp—Reply using an IPv4 or IPv6 UDP packet.

no-reply—Do not reply to the ping request.

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).

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 l2vpn instance on page 3136](#)

[ping mpls l2vpn instance detail on page 3136](#)

[ping mpls l2vpn interface <interface-name> reply-mode on page 3136](#)

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 l2vpn instance

```
user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2
```

```
!!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l2vpn instance detail

```
user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2 detail
```

```
Request for seq 1, to interface 68, labels <800001, 100176>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <800001, 100176>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <800001, 100176>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <800001, 100176>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <800001, 100176>
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l2vpn interface <interface-name> reply-mode

```
user@host> ping mpls l2vpn interface lt-1/2/0.21 reply-mode ip-udp
```

```
!!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


ping mpls l3vpn

Syntax

```
ping mpls l3vpn prefix prefix-name
<l3vpn-name>
<bottom-label-ttl>
<count count>
<destination address>
<detail>
<exp forwarding-class>
<logical-system (all | logical-system-name)>
<size bytes>
<source source-address>
<sweep>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

The **size** and **sweep** options were introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a **ping mpls l3vpn** command.

Options

bottom-label-ttl—(Optional) Display the time-to-live value for the bottom label in the label stack.

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.

l3vpn-name—(Optional) Layer 3 VPN name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

prefix *prefix-name*—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.

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).

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 3138](#)

[ping mpls l3vpn detail on page 3139](#)

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
```


request mpls container-lsp

Syntax

```
request mpls container-lsp
<logical-system (all | logical-system-name)>
<name lsp-name>
<adjust-autobandwidth>
<normalization>
```

Release Information

Command introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

Description

Manually trigger a bandwidth allocation adjustment for the container label-switched path (LSP).

Options

none—Manually trigger a bandwidth allocation adjustment for all active member LSP paths.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name *lsp-name*—(Optional) Manually trigger a bandwidth allocation adjustment on the specified member LSP only.

adjust-autobandwidth—(Optional) Request LSP autobandwidth adjustment.

normalization—(Optional) Request container LSP normalization.

Required Privilege Level

clear, maintenance

RELATED DOCUMENTATION

[show mpls container-lsp | 3179](#)

[clear mpls container-lsp | 3105](#)

List of Sample Output

[request mpls container-lsp on page 3141](#)

[request mpls container-lsp on page 3141](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request mpls container-lsp

user@host> request mpls container-lsp *lsp-name* normalize

request mpls container-lsp

user@host> request mpls container-lsp normalize bandwidth *bps*

request mpls lsp adjust-autobandwidth

List of Syntax

[Syntax on page 3142](#)

[Syntax \(EX and QFX Series Switches\) on page 3142](#)

Syntax

```
request mpls lsp adjust-autobandwidth
<logical-system (all | logical-system-name)>
<name lsp-name>
```

Syntax (EX and QFX Series Switches)

```
request mpls lsp adjust-autobandwidth
<name lsp-name>
```

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.5 for EX Series switches.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

Description

Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).

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.

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 *make-before-break*.

Options

none—Manually trigger a bandwidth allocation adjustment for all active LSP paths.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name *lsp-name*—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.

Additional Information

For this command to work properly, the following conditions must exist:

- Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the **request mpls lsp adjust-autobandwidth** command.
- The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.

Required Privilege Level

clear, maintenance

RELATED DOCUMENTATION

[auto-bandwidth | 2559](#)

[Configuring Automatic Bandwidth Allocation for LSPs | 547](#)

List of Sample Output

[request mpls lsp adjust-auto-bandwidth on page 3143](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
request mpls lsp adjust-auto-bandwidth
```

```
user@host> request mpls lsp adjust-auto-bandwidth
```


show connections

List of Syntax

[Syntax on page 3144](#)

[Syntax \(EX Series Switches\) on page 3144](#)

Syntax

```
show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-interface-switch>
<down | up | up-down>
<history>
<labels>
<logical-system (all | logical-system-name)>
<name>
<status>
```

Syntax (EX Series Switches)

```
show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-interface-switch>
<down | up | up-down>
<history>
<labels>
<name>
<status>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display information about the configured circuit cross-connect (CCC) connections.

Options

none—Display the standard level of output for all configured CCC connections.

all—(Optional) Display all connections.

brief | extensive—(Optional) Display the specified level of output. Use history to display information about connection history. Use labels to display labels used for transmit and receive LSPs. Use status to display information about the connection and interface status.

interface-switch—(Optional) Display interface switch connections only.

lsp-switch—(Optional) Display LSP switch connections only.

p2mp-receive-switch—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.

p2mp-transmit-switch—(Optional) Display local interface to point-to-multipoint LSP switch connections only.

remote-interface-switch—(Optional) Display remote interface switch connections only.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

labels—(Optional) Display labels used for transmit and receive.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display information about the specified connection only.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level

view

Output Fields

[Table 83 on page 3145](#) describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 83: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On Off]	Whether link monitoring is enabled: On or Off .
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.

Table 83: show connections Output Fields (*continued*)

Field Name	Field Description
Legend for connection types	Type of connection: <ul style="list-style-type: none"> • if-sw—Layer 2 switching cross-connect. • rmt-if—Remote interface switch. While graceful restart is in progress, rmt-if will display a state (St) of Restart. • lsp-sw—LSP stitching cross-connect. While graceful restart is in progress, lsp-sw will display a state (St) of Restart.
Legend for circuit types	Type of circuits: <ul style="list-style-type: none"> • intf—Interface circuit. • tlsp—Transmit LSP circuit. • rlsp—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the Up (operational) state.
# Up trans	Number of times that the connection or circuit has transitioned to the Up (operational) state.

Sample Output

show connections

user@switch> **show connections**

```
CCC and TCC connections [Link Monitoring On]
Legend for status (St)           Legend for connection types
UN -- uninitialized             if-sw:  interface switching
NP -- not present               rmt-if: remote interface switching
WE -- wrong encapsulation       lsp-sw: LSP switching
DS -- disabled
Dn -- down                     Legend for circuit types
-> -- only outbound conn is up  intf -- interface
```



```

<- -- only inbound conn is up      tlsp -- transmit LSP
Up -- operational                    rlsp -- receive LSP
RmtDn -- remote CCC down
Restart -- restarting

```

CCC Graceful restart : Restarting

Connection/Circuit	Type	St	Time last up	# Up trans
IFSW-ed	if-sw	Up	Aug 5 15:39:15	1
so-1/0/2.0	intf	Up		
t1-0/1/2.0	intf	Up		
SW-db	rmt-if	Restart		0
so-1/0/3.0	intf	Up		
pro4-ca	tlsp	Dn		
pro4-ac	rlsp	NP		

show dynamic-tunnels database

Syntax

```
show dynamic-tunnels database
<destination>
<logical-system (all | logical-system-name) >
<table routing-table-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display dynamic tunnel database information.

Options

none—Display dynamic tunnel database information for all destinations and routing tables.

destination—(Optional) Display database entries for the specified IP address (with optional destination prefix length) only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display database entries for the specified table only.

Required Privilege Level

view

List of Sample Output

[show dynamic-tunnels database \(Tunnel Is Up\) on page 3150](#)

[show dynamic-tunnels database \(No Tunnel PIC\) on page 3150](#)

[show dynamic-tunnels database \(Tunnel Is Expiring\) on page 3150](#)

[show dynamic-tunnels database \(Destination Specified\) on page 3151](#)

[show dynamic-tunnels database \(Localization\) on page 3151](#)

[show dynamic-tunnels database \(MPLS-over-UDP Dynamic Tunnels on PTX Series Routers and QFX Series Switches\) on page 3151](#)

[show dynamic-tunnels database \(Segment Routing LSPs\) on page 3152](#)

Output Fields

[Table 84 on page 3149](#) lists the output fields for the **show dynamic-tunnels database** command. Output fields are listed in the approximate order in which they appear.

Table 84: show dynamic-tunnels database Output Fields

Field Name	Field Description
Table	Name of the routing table (for example, inet.0).
Destination-network	Destination IP address and subnet.
Tunnel to	Destination IP address and prefix of the tunnel.
State	State of the tunnel: Up , Up (expires in <i>nn:nn:nn</i>seconds) , or Dn (down).
Reference count	Number of routes across the dynamic tunnel that are currently being resolved.
Next-hop type	Type of tunnel: GRE or UDP (BGP-Signal). <ul style="list-style-type: none"> • GRE or UDP (BGP signal) • SRTE—Segment routing traffic-engineered LSP.
Source address	Source IP address of the tunnel.
Next-hop	IP address of the destination interface.
VPN Label	The label provided by the peer device to identify the VPN through which the packet needs to go. This label is used to identify the VRF for route lookup.
Ingress Route	The IGP route along with the corresponding metric that has been selected for forwarding the tunnel-encapsulated packet.
Localized PFE	Packet Forwarding Engine interface which is the anchor Packet Forwarding Engine for the localized next-hop-based dynamic tunnels. When the anchor Packet Forwarding Engine of the tunnel goes down, it is represented by a # near the Packet Forwarding Engine name.
LSP template name	Name of the segment routing traffic-engineered template configured for dynamic creation of segment routing LSPs.
State	State of the destination interface: Up , Dn , or Dn (no tunnel pic) .
Status	Status of the dynamic segment routing LSP.

Sample Output

show dynamic-tunnels database (Tunnel Is Up)

```
user@host> show dynamic-tunnels database
```

```
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32
Reference count: 4
Next-hop type: UDP
Source address: 10.255.120.92
Next hop: tunnel-composite, 0x31132f64, nhid 3406
VPN Label: Push 120 Reference count: 3
Ingress Route: 10.255.120.94/32, via metric 2
Traffic Statistics: Packets 241367951, Bytes 356741831578
State: Up
```

show dynamic-tunnels database (No Tunnel PIC)

```
user@host> show dynamic-tunnels database
```

```
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Dn
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
State: Dn (no tunnel pic)
```

show dynamic-tunnels database (Tunnel Is Expiring)

```
user@host> show dynamic-tunnels database
```

```
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up (expires in 00:14:56 seconds)
Reference count: 0
```



```

Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up

```

show dynamic-tunnels database (Destination Specified)

```
user@host> show dynamic-tunnels database 10.255.120.94
```

```

Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up

```

show dynamic-tunnels database (Localization)

```
user@host> show dynamic-tunnels database
```

```

Destination-network: 1.0.0.0/8
Tunnel to: 1.1.1.6/32
Reference count: 5
Next-hop type: UDP
Source address: 1.1.1.2
Next hop: tunnel-composite, 0xc807930, nhid 1016
Localized PFE: pfe-1/0/0
VPN Label: Push 299808 Reference count: 4
Ingress Route: 1.1.1.6/32, via metric 2
Traffic Statistics: Packets 0, Bytes 0
State: Up

```

show dynamic-tunnels database (MPLS-over-UDP Dynamic Tunnels on PTX Series Routers and QFX Series Switches)

```
user@host> show dynamic-tunnels database
```

```

*- Signal Tunnels #- PFE-down
Table: inet.3

```



```
Destination-network: 22.33.0.0/16
```

```
Destination-network: 22.33.44.0/24
```

show dynamic-tunnels database (Segment Routing LSPs)

```
user@host> show dynamic-tunnels database
```

```
Table: inetcolor.0
```

```
Destination-network: 22.33.44.0:0/24
```

```
Tunnel to: 22.33.44.55:124/64
```

```
Reference count: 2
```

```
Next-hop type: SRTE
```

```
LSP template name: 22.33.44.55:7c:dt-srte-tunnel1
```

```
Status: Initiated/Established
```


show link-management

Syntax

```
show link-management
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.

Options

This command has no options.

Required Privilege Level

view

RELATED DOCUMENTATION

show link-management peer 3157
show link-management routing 3160
show link-management statistics 3164
show link-management te-link 3167

List of Sample Output

[show link-management on page 3155](#)

Output Fields

[Table 85 on page 3153](#) describes the output fields for the **show link-management** command. Output fields are listed in the approximate order in which they appear.

Table 85: show link-management Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down.

Table 85: show link-management Output Fields (*continued*)

Field Name	Field Description
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296 .
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296 .
State	State of the control channel: Up or Down .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295 .
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295 .
Flags	Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.
TE link name	Name of the traffic-engineered link.
State	State of the traffic-engineered link: Up , Down , or Init .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH .
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet .

Table 85: show link-management Output Fields (*continued*)

Field Name	Field Description
Minimum bandwidth	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
Maximum bandwidth	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
Name	Name of the interface.
State	State of the interface: Up or Down .
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: Yes or No .
LSP-name	LSP name.

Sample Output

show link-management

user@host> **show link-management**

```

Peer name: PEER-A, System identifier: 11973
State: Up, Control address: 10.255.245.4
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
      24547      24547 Up          1027      1026
TE links:
  pro4-ba

```


TE link name: pro4-ba, State: Init
Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
PSC-1,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps,
Available bandwidth: 155.52Mbps

Name	State	Local ID	Remote ID	Bandwidth Used	LSP-name
so-1/0/2	Up	21271	0	155.52Mbps	No

show link-management peer

Syntax

```
show link-management peer
<name peer-name>
```

Release Information

Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display Multiprotocol Label Switching (MPLS) peer link information.

Options

- none**—Display all peer link information.
- name *peer-name***—(Optional) Display information for the specified peer only.

Required Privilege Level

view

RELATED DOCUMENTATION

show link-management 3153
show link-management routing 3160
show link-management statistics 3164
show link-management te-link 3167

List of Sample Output

[show link-management peer on page 3158](#)

Output Fields

[Table 86 on page 3157](#) describes the output fields for the **show link-management peer** command. Output fields are listed in the approximate order in which they appear.

Table 86: 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.

Table 86: show link-management peer Output Fields (*continued*)

Field Name	Field Description
State	State of the peer: Up or Down .
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296 .
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296 .
State	State of the control channel: Up or Down .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295 .
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295 .
Flags	Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.

Sample Output

show link-management peer

user@host> **show link-management peer**

```
Peer name: sonet, System identifier: 41448
  State: Up, Control address: 70.70.70.70
  Hello interval: 10000, Hello dead interval: 30000
```


CC local ID	CC remote ID	State	TxSeqNum	RcvSeqNum	Flags
3265	0	ConfSnd	1	0	R
TE links:					
to-sonet					

show link-management routing

Syntax

```
show link-management routing
<peer <name name> | te-link <name name>>
<resource <name name>>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.

Options

- none**—Display all peer and traffic-engineered link information.
- peer <name name>**—(Optional) Display information for all peers or for the specified peer only.
- resource <name name>**—(Optional) Display information for all resources or for the specified resource only.
- te-link <name name>**—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.

Required Privilege Level

view

RELATED DOCUMENTATION

show link-management 3153
show link-management peer 3157
show link-management statistics 3164
show link-management te-link 3167

List of Sample Output

[show link-management routing on page 3162](#)

Output Fields

[Table 87 on page 3161](#) describes the output fields for the **show link-management routing** command. Output fields are listed in the approximate order in which they appear.

Table 87: 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.
State	State of the control channel.
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: Up or Down .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Resource	Forwarding adjacency LSP information.

Table 87: show link-management routing Output Fields (*continued*)

Field Name	Field Description
Type	Type of resource. The type is always a forwarding adjacency LSP.
State	State of the LSP: Up or Down .
System Identifier	Internal identifier for the peer. The range of values is 0 through 64,000 .
Total bandwidth	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.
Traffic parameters	<ul style="list-style-type: none"> • Encoding—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET, Ethernet, and Packet. • Switching—Type of switching that can be performed on the traffic-engineered link: PSC-1 and Packet. • Granularity—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always unknown.

Sample Output

show link-management routing

user@host> **show link-management routing**

```

Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
State: Up, Control address: (null)
  Control-channel          State
  fe-0/1/0.0              Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
State: Up, Control address: (null)
  Control-channel          State
  fe-0/1/2.0              Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
State: Down, Control address: (null)
  Control-channel          State
  so-0/2/0.0              State

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
State: Down, Control address: (null)
  Control-channel          State

```


so-0/2/1.0

...

TE link name: __rpd:fe-0/1/0.0, State: Up
 Local identifier: 2147483649, Remote identifier: 0,
 Local address: 192.168.37.66, Remote address: 192.168.37.66,
 Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
 Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:fe-0/1/2.0, State: Up
 Local identifier: 2147483650, Remote identifier: 0,
 Local address: 192.168.37.73, Remote address: 192.168.37.73,
 Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
 Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:so-0/2/0.0, State: Down
 Local identifier: 2147483651, Remote identifier: 0,
 Local address: 192.168.37.82, Remote address: 192.168.37.95,
 Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
 Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
 Total bandwidth: 0bps, Traffic parameters: Encoding: Packet, Switching: Packet,
 Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up System identifier: 2147483654,
 Total bandwidth: bw[1]=10Mbps, Traffic parameters: Encoding: Packet,
 Switching: Packet, Granularity: Unknown

show link-management statistics

Syntax

```
show link-management statistics
<peer <name name>>
```

Release Information

Command introduced in Junos OS Release 8.0.
 Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display statistical information for Link Management Protocol (LMP) packets.

Options

- none**—Display information for all peers.
- peer <name name>**—(Optional) Display information for all peers or for the specified peer only.

Required Privilege Level

view

RELATED DOCUMENTATION

show link-management	 3153
show link-management peer	 3157
show link-management routing	 3160
show link-management te-link	 3167

List of Sample Output

[show link-management statistics on page 3165](#)

Output Fields

[Table 88 on page 3164](#) describes the output fields for the **show link-management statistics** command. Output fields are listed in the approximate order in which they appear.

Table 88: 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.

Table 88: show link-management statistics Output Fields (*continued*)

Field Name	Field Description
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.
Stale acknowledgments	Number of configAck and LinkSummaryAck packets received that have a stale message ID.
Stale negative acknowledgments	Number of configNack and LinkSummaryNack packets received that have a stale message ID.
Sent packets	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Retransmitted packets	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Dropped packets	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

Sample Output

```
show link-management statistics
```

```
user@host> show link-management statistics peer pro4-a
```



```
Statistics for peer pro4-a
  Received packets
    Config: 1
    Hello: 2572
  Small packets: 0
  Wrong protocol version: 0
  Messages for unknown peer: 0
  Messages for bad state: 0
  Stale acknowledgments: 0
  Stale negative acknowledgments: 0
  Sent packets
    Config: 2
    ConfigAck: 1
    Hello: 2572
  Retransmitted packets
    Config: 1
```


show link-management te-link

Syntax

```
show link-management te-link
<brief | detail>
<name name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

Options

none—Display information for all traffic-engineered links.

brief | detail—(Optional) Display the specified level of output.

name *name*—(Optional) Display information for the specified traffic-engineered link only.

Required Privilege Level

view

RELATED DOCUMENTATION

[show link-management](#) | 3153

[show link-management peer](#) | 3157

[show link-management routing](#) | 3160

[show link-management statistics](#) | 3164

List of Sample Output

[show link-management te-link on page 3169](#)

Output Fields

[Table 89 on page 3168](#) describes the output fields for the **show link-management te-link** command. Output fields are listed in the approximate order in which they appear.

Table 89: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: Up or Down .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH .
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
Available Bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Name	Name of the interface.
State	State of the interface: Up or Down .
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: Yes or No .

Table 89: show link-management te-link Output Fields (continued)

Field Name	Field Description
LSP-name	LSP name.

Sample Output

show link-management te-link

user@host> show link-management te-link

```

TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name          State Local ID Remote ID      Bandwidth Used LSP-name
    falsp-bd      Dn          43077         0           0bps No

TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
  Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
  Available bandwidth: 8Mbps
    Name          State Local ID Remote ID      Bandwidth Used LSP-name
    falsp-be      Up          43076         0          10Mbps Yes  e2elsp-bf

```


show mpls abstract-hop-membership

Syntax

```
show mpls abstract-hop-membership
<abstract-hop-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 17.1 for all platforms.

Description

Display MPLS abstract hop membership tables for each abstract hop configured on the device.

Options

none—(Optional) Display the MPLS abstract hop membership table for all the configured abstract hops on the router.

abstract-hop-name—(Optional) Display the MPLS abstract hop membership table for the specified abstract hop.

instance instance-name—(Optional) Display the MPLS abstract hop membership table for the specified instance. If *instance-name* is omitted, information is displayed for the master instance.

logical-system (all | logical-system-name)—(Optional) Display the MPLS abstract hop membership table for all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[Example: Configuring Abstract Hops for MPLS LSPs | 446](#)

[abstract-hop | 2534](#)

[constituent-list | 2578](#)

[show mpls lsp abstract-computation | 3240](#)

List of Sample Output

[show mpls abstract-hop-membership on page 3171](#)

Output Fields

Table 90 on page 3171 describes the output fields for the **show mpls abstract-hop-membership** command. Output fields are listed in the approximate order in which they appear.

Table 90: show mpls abstract-hop-membership Output Fields

Field Name	Field Description
Abstract hop	Name of the abstract hop.
Credibility	Credibility value associated with the interior gateway protocol in use.
Address	IP address of the abstract hop member nodes.

Sample Output

show mpls abstract-hop-membership

user@host> **show mpls abstract-hop-membership**

Abstract hop: ah1

Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.1
Address: 127.0.0.2
Address: 127.0.0.3

Abstract hop: ah2

Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.3
Address: 127.0.0.4

Abstract hop: ah3

Credibility: 0
Address: 127.0.0.6
Address: 127.0.0.3
Address: 127.0.0.5

show mpls admin-groups

List of Syntax

[Syntax on page 3172](#)

[Syntax \(EX Series Switches\) on page 3172](#)

Syntax

```
show mpls admin-groups
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls admin-groups
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display information about configured Multiprotocol Label Switching (MPLS) administrative groups.

Options

none—Display information about the configured MPLS administrative groups.

instance *instance-name*—(Optional) Display MPLS administrative group information for the specified instance. If ***instance-name*** is omitted, MPLS administrative group information 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.

Required Privilege Level

view

List of Sample Output

[show mpls admin-groups on page 3173](#)

Output Fields

[Table 91 on page 3173](#) describes the output fields for the **show mpls admin-groups** command. Output fields are listed in the approximate order in which they appear.

Table 91: show mpls admin-groups Output Fields

Field Name	Field Description
Group	Name of the administrative group.
Bit index	Value assigned to the administrative group.

Sample Output

show mpls admin-groups

user@host> **show mpls admin-groups**

Group	Bit index
black	3
blue	2
gold	1
green	0

show mpls association

Syntax

```
show mpls association (iif incoming-interface | oif outgoing-interface)
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 16.1 for the M Series, MX Series, and T Series.

Description

Display the Multiprotocol Label Switching (MPLS) label-switched paths (LSPs) based on the association with an incoming or outgoing LSP interface. The command output displays the list of RSVP-TE LSPs carrying traffic in and out of the same interface.

Options

iif *incoming-interface-name*—Display list of RSVP-TE LSPs that share the specified incoming interface to bring in traffic. This option works on transit label-switching routers (LSRs) and egress label edge routers (LERs).

oif *outgoing-interface-name*—Display list of RSVP-TE LSPs that share the specified outgoing interface to carry out traffic. This option works on ingress LERs and transit LSRs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[show mpls correlation nexthop-id](#) | 3192

List of Sample Output

[show mpls association iif on page 3175](#)

[show mpls association oif on page 3175](#)

Output Fields

[Table 92 on page 3175](#) describes the output fields for the **show mpls association** command. Output fields are listed in the approximate order in which they appear.

Table 92: show mpls association Output Fields

Field Name	Field Description
To	Destination IP address of the corresponding LSP.
From	Source IP address of the corresponding LSP.
State	State of the corresponding LSP handled by this RSVP session: Up , Dn (down), or Restart .
LSPname	Name of the LSP.

Sample Output

show mpls association iif

```
user@host> show mpls association iif ge-0/0/0.0
```

To	From	State	LSPname
128.102.174.121	128.102.180.21	Up	LSP-ABC
128.102.174.121	128.102.180.21	Up	LSP-ABC1
128.102.174.121	128.102.180.21	Up	LSP-ABC2
128.102.174.121	128.102.180.21	Up	LSP-ABC3
Total 4 displayed, Up 4, Down 0			

show mpls association oif

```
user@host> show mpls association oif ge-0/0/0.0
```

To	From	State	LSPname
128.102.174.121	128.102.180.21	Up	LSP-ABC
128.102.174.121	128.102.180.21	Up	LSP-ABC1
128.102.174.121	128.102.180.21	Up	LSP-ABC2
128.102.174.121	128.102.180.21	Up	LSP-ABC3

show mpls call-admission-control

List of Syntax

[Syntax on page 3176](#)

[Syntax \(EX Series Switches\) on page 3176](#)

Syntax

```
show mpls call-admission-control
<instance instance-name>
<logical-system (all | logical-system-name)>
<lsp-name>
```

Syntax (EX Series Switches)

```
show mpls call-admission-control
<lsp-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.

Options

none—Display CAC information for all LSPs.

instance *instance-name*—(Optional) Display MPLS LSP CAC information for the specified instance. If *instance-name* is omitted, MPLS LSP CAC information for the master instance is displayed.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-name—(Optional) Display CAC information for the specified LSP only.

Additional Information

The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.

Required Privilege Level

view

List of Sample Output

[show mpls call-admission-control on page 3177](#)

Output Fields

[Table 93 on page 3177](#) describes the output fields for the **show mpls call-admission-control** command.

Output fields are listed in the approximate order in which they appear.

Table 93: show mpls call-admission-control Output Fields

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at ct0) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.
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: prol-be
*Primary
  Available bandwidth: 0bps
```


LSP name: prol-be-1

*Primary

Available bandwidth: 60kbps

LSP name: prol-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: prol-be-gold-2

*Primary

Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: prol-be-silver

*Primary prim1

Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>

Layer2 connections:

Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)

VC bandwidth: <ct0 20kbps> <ct1 20kbps>

Standby sec1

Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>

Layer2 connections:

Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)

VC bandwidth: <ct0 20kbps> <ct1 20kbps>

show mpls container-lsp

Syntax

```
show mpls container-lsp
<brief | detail | extensive | terse>
<count-active-routes>
<defaults>
<descriptions>
<down | up>
<egress>
<ingress>
<logical-system (all | logical-system-name)>
<name name>
<statistics>
<transit>
<unidirectional>
```

Release Information

Command introduced in Junos OS Release 14.2.

Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.

Description

Display information about configured and active Multiprotocol Label Switching (MPLS) container label-switched paths (LSPs).

Options

none—Display standard information about all configured and active member LSPs of the container LSP.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

count-active-routes—(Optional) Show active routes for the container LSP.

defaults—(Optional) Display the default settings of the container LSP.

descriptions—(Optional) Display the container LSP descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls container-lsp]** hierarchy level. Only the LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

egress—(Optional) Display the LSPs ending at this device.

NOTE: The **egress** option displays all the LSPs including regular LSPs, members of container LSPs, and transit LSPs. This is an expected behavior for all platforms.

- ingress**—(Optional) Display the member LSPs originating from this device.
- logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.
- name *name***—(Optional) Display information about the specified LSP or group of LSPs.
- statistics**—(Optional) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.
- transit**—(Optional) Display LSPs transiting this routing device.
- unidirectional**—(Optional) Display unidirectional LSP information.

Required Privilege Level
view

RELATED DOCUMENTATION

- [request mpls container-lsp | 3140](#)
- [clear mpls container-lsp | 3105](#)

List of Sample Output
[show mpls container-lsp on page 3185](#)
[show mpls container-lsp extensive on page 3185](#)

Output Fields

[Table 94 on page 3180](#) describes the output fields for the **show mpls container-lsp** command. Output fields are listed in the approximate order in which they appear.

Table 94: show mpls container-lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSP	Information about the member LSPs on the ingress routing device. Each LSP has one line of output.	All levels

Table 94: show mpls container-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Container LSP name	Name of the container LSP.	All levels
Member LSP count	Number of member LSPs in the container LSP.	All levels
To	Destination (egress routing device) of the session.	brief
From	Source (ingress routing device) of the session.	brief detail
State	State of the LSP handled by this RSVP session: <ul style="list-style-type: none"> • Up • Dn (down) • Restart 	brief detail
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	brief
ActivePath	(Ingress LSP) Name of the active path: Primary or Secondary.	detail extensive
LSPname	Name of the member LSP.	brief detail
Egress LSP	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
Transit LSP	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
Min LSPs	Minimum number of member LSPs. Default: 1	extensive
Max LSPs	Number of member LSPs that the container LSP can have at maximum. Default: 64 (due to ECMP limit)	extensive

Table 94: show mpls container-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Aggregate bandwidth	Sum of the bandwidths of all member LSPs.	extensive
NormalizeTimer	Duration between two normalization events. When not configured, 21600 seconds (6 hours) is set as the default value.	extensive
NormalizeThreshold	Change in aggregate LSP utilization to trigger splitting or merging expressed in percentage.	extensive
Max Signaling BW	Maximum bandwidth used to signal LSPs after a normalization event. Default value is 0 bps. When not configured, the value is inherited from the splitting bandwidth configuration. NOTE: Between two normalization events, when auto-bandwidth adjustment happens, the per-LSP auto-bandwidth configuration and thresholds are used, instead of the maximum signaling bandwidth threshold.	extensive
Min Signaling BW	Minimum bandwidth used to signal LSPs after a normalization event. Default value is 0 bps. When not configured, the value is inherited from the merging bandwidth configuration. NOTE: Between two normalization events, when auto-bandwidth adjustment happens, the per-LSP auto-bandwidth configuration and thresholds are used, instead of the minimum signaling bandwidth threshold.	extensive
Splitting BW	Bandwidth used for LSP splitting and merging. Default value is 0 bps. When not configured, the value is inherited from the auto-bandwidth maximum bandwidth configuration.	extensive
Merging BW	Bandwidth used for LSP splitting and merging. Default value is 0 bps. When not configured, the value is inherited from the auto-bandwidth minimum bandwidth configuration.	extensive
LSPtype		extensive
LoadBalance		extensive
MinBW	Minimum LSP bandwidth in bps related to auto-bandwidth.	extensive

Table 94: show mpls container-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
AdjustTimer	Total amount of time in seconds allowed before LSP bandwidth adjustment take place. Range: 300 through 315360000 seconds	extensive
Max AvgBW util	Current value of the actual maximum average bandwidth utilization in bps.	extensive
Overflow limit	Threshold overflow limit.	extensive
Underflow limit	Threshold underflow limit.	extensive
Encoding type		extensive
Switching type		extensive
GPID		extensive
Priorities	Setup priority and hold priority values. For setup priority, 0 and 7 is the highest and lowest priority, respectively. When not explicitly configured, 7 and 0 are set as the default values for the setup priority and hold priority, respectively.	extensive
Bandwidth		extensive
SmartOptimizeTimer	Time in seconds allowed before path reoptimization.	extensive
Computed ERO	Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	extensive

Table 94: show mpls container-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Received RRO	<p>Received record route.</p> <p>RRO is a series of hops, each with an address followed by a flag. In most cases, the received RRO is the same as the computed ERO. If the received RRO is different from the computed ERO, there is a topology change in the network, and the route is taking a detour.</p> <p>The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> • 0x01—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding path message. • 0x02—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • 0x03—Combination of 0x01 and 0x02. • 0x04—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • 0x08—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the local protection available bit is set but the node protection bit is cleared. • 0x09—Detour is established. Combination of 0x01 and 0x08. • 0x10—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • 0x20—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently. • 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08. 	extensive
Make-before-break		extensive
Record Route		extensive

Table 94: show mpls container-lsp Output Fields (continued)

Field Name	Field Description	Level of Output
Automatic Autobw adjustment succeeded		extensive
CSPF		extensive
Created	Date and time the LSP was created.	extensive

Sample Output

show mpls container-lsp

```
user@host> show mpls container-lsp
```

```
Ingress LSP: 1 sessions
Container LSP name          Member LSP count
test                        2
To          From          State Rt P    ActivePath    LSPname
10.255.107.76 10.255.107.78 Up    0 *          test-1
10.255.107.76 10.255.107.78 Up    0 *          test-2
Total 2 displayed, Up 2, Down 0

naling Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls container-lsp extensive

```
user@host> show mpls container-lsp extensive
```

```
Ingress LSP: 1 sessions
Container LSP name: test, Member count: 2
Normalization
  Min LSPs: 2, Max LSPs: 64, Aggregate bandwidth: 0bps
  NormalizeTimer: 1800 secs, NormalizeThreshold: 0%
```


Max Signaling BW: 2kbps, Min Signaling BW: 2kbps, Splitting BW: 5Mbps, Merging BW: 2kbps

Normalization in 989 second(s)

10.255.107.76

From: 10.255.107.78, State: Up, ActiveRoute: 0, LSPname: test-1

ActivePath: (primary)

LSPtype: Dynamic Configured, Penultimate hop popping

LoadBalance: Random

Autobandwidth

MinBW: 1000bps

AdjustTimer: 300 secs

Max AvgBW util: 0bps, Bandwidth Adjustment in 89 second(s).

Overflow limit: 0, Overflow sample count: 0

Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps

Encoding type: Packet, Switching type: Packet, GPID: IPv4

*Primary State: Up, No-decrement-ttl

Priorities: 7 0

Bandwidth: 1000bps

SmartOptimizeTimer: 180

Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)

1.3.0.2 S 1.7.0.1 S

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):

1.3.0.2 1.7.0.1

11 Jul 13 20:08:26.613 Make-before-break: Switched to new instance

10 Jul 13 20:08:04.360 Record Route: 1.3.0.2 1.7.0.1

9 Jul 13 20:08:04.360 Up

8 Jul 13 20:08:04.360 Automatic Autobw adjustment succeeded: BW changes from 2000 bps to 1000 bps

7 Jul 13 20:08:04.314 Originate make-before-break call

6 Jul 13 20:08:04.314 CSPF: computation result accepted 1.3.0.2 1.7.0.1

5 Jul 13 20:05:02.423 Selected as active path

4 Jul 13 20:05:02.422 Record Route: 1.3.0.2 1.7.0.1

3 Jul 13 20:05:02.421 Up

2 Jul 13 20:05:02.376 Originate Call

1 Jul 13 20:05:02.376 CSPF: computation result accepted 1.3.0.2 1.7.0.1

Created: Sat Jul 13 20:03:03 2013

10.255.107.76

From: 10.255.107.78, State: Up, ActiveRoute: 0, LSPname: test-2

ActivePath: (primary)

LSPtype: Dynamic Configured, Penultimate hop popping

LoadBalance: Random

Autobandwidth

MinBW: 1000bps


```

AdjustTimer: 300 secs
Max AvgBW util: 0bps, Bandwidth Adjustment in 89 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up, No-decrement-ttl
  Priorities: 7 0
  Bandwidth: 1000bps
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
1.2.0.2 S 1.4.0.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    1.2.0.2 1.4.0.2
11 Jul 13 20:08:05.363 Make-before-break: Switched to new instance
10 Jul 13 20:08:04.450 Record Route:  1.2.0.2 1.4.0.2
 9 Jul 13 20:08:04.449 Up
 8 Jul 13 20:08:04.449 Automatic Autobw adjustment succeeded: BW changes from
2000 bps to 1000 bps
 7 Jul 13 20:08:04.327 Originate make-before-break call
 6 Jul 13 20:08:04.327 CSPF: computation result accepted  1.2.0.2 1.4.0.2
 5 Jul 13 20:05:00.849 Selected as active path
 4 Jul 13 20:05:00.841 Record Route:  1.3.0.2 1.7.0.1
 3 Jul 13 20:05:00.831 Up
 2 Jul 13 20:05:00.513 Originate Call
 1 Jul 13 20:05:00.502 CSPF: computation result accepted  1.3.0.2 1.7.0.1
Created: Sat Jul 13 20:03:03 2013
Total 2 displayed, Up 2, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```


show mpls context-identifier

Syntax

```
show mpls context-identifier
<brief | detail>
<logical-system (all | logical-system-name)>
<primary>;
<protector>;
```

Release Information

Command introduced in Junos OS Release 11.4R3.

Description

Display information about configured egress protection context identifiers.

Options

none—Display standard information about egress protection.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

primary—(Optional) Perform this operation on the primary node.

protector—(Optional) Perform this operation on the protector node.

Required Privilege Level

view

RELATED DOCUMENTATION

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

Example: Configuring MPLS Egress Protection for Layer 3 VPN Services

List of Sample Output

[show mpls context-identifier detail \(Protector\) on page 3189](#)

[show mpls context-identifier detail \(Primary\) on page 3190](#)

Output Fields

[Table 95 on page 3189](#) describes the output fields for the **show mpls egress-protection detail** command. Output fields are listed in the approximate order in which they appear.

Table 95: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
ID	Context identifier.	All levels
Type	Indicates node type: protector or primary	All levels
Metric	MPLS cost value of the context identifier route. This route appears in inet.0 on the protector and primary nodes. On the protector node, the metric is a larger number.	All levels
Mode	Indicates <code>advertise-mode</code> : proxy or alias	detail
Context table	Name of the MPLS routing table created for egress protection.	All levels
Context LSPs	Names of the LSPs that have egress protection configured. Loopback interface addresses of the devices from which the LSPs are originated.	detail
Total	Total number of primary and protector nodes.	All levels
Primary	Number of primary nodes.	All levels
Protector	Number of protector nodes.	All levels

Sample Output

show mpls context-identifier detail (Protector)

```
user@host> show mpls context-identifier detail
```

```
ID: 166.1.3.1
  Type: protector, Metric: 16777215, Mode: alias
  Context table: __166.1.3.1__.mpls.0, Label out: 299968
```


Sample Output

show mpls context-identifier detail (Primary)

user@host> **show mpls context-identifier detail**

```
ID: 166.1.3.1
```

```
  Type: primary, Metric: 1, Mode: alias
```

```
Total 1, Primary 1, Protector 0
```


show mpls correlation label

Syntax

```
show mpls correlation label label-value  
<brief | detail | extensive | terse>  
<descriptions>  
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 15.2 for the M Series, MX Series, and T Series.

Description

Display the correlation information for the Multiprotocol Label Switching (MPLS) label-switched path (LSP) with the owner of the label.

Options

label-value—Display information about the specified label.

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.

descriptions—(Optional) Display the LSP descriptions. To view this information, you must configure the description statement at the [**edit protocol mpls lsp**] hierarchy level. Only the LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[show mpls correlation nexthop-id](#) | 3192

[show mpls association](#) | 3174

show mpls correlation nexthop-id

Syntax

```
show mpls correlation nexthop-id nexthop-id
<descriptions>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 16.1 for the M Series, MX Series, and T Series.

Description

Display the correlation information for the Multiprotocol Label Switching (MPLS) label-switched path (LSP) with the owner of the next-hop ID.

Options

nexthop-id—Display information about the specified next-hop ID.

descriptions—(Optional) Display the LSP descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls lsp]** hierarchy level. Only the LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[show mpls association](#) | [3174](#)

List of Sample Output

[show mpls correlation nexthop-id on page 3193](#)

Output Fields

[Table 96 on page 3193](#) describes the output fields for the **show mpls correlation nexthop-id** command. Output fields are listed in the approximate order in which they appear.

Table 96: show mpls correlation nexthop-id Output Fields

Field Name	Field Description
LSP name	Name of the LSP associated with the specified next-hop ID.

Sample Output

show mpls correlation nexthop-id

user@host> **show mpls correlation nexthop-id *nexthop-id***

```
LSP name: LSP-ABC
```


show mpls cspf

List of Syntax

[Syntax on page 3194](#)

[Syntax \(EX Series Switches\) on page 3194](#)

Syntax

```
show mpls cspf
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls cspf
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.

Options

none—Display MPLS CSFP statistics.

instance *instance-name*—(Optional) Display MPLS CSPF information for the specified instance. If *instance-name* is omitted, MPLS CSPF information 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.

Required Privilege Level

view

List of Sample Output

[show mpls cspf on page 3196](#)

Output Fields

[Table 97 on page 3195](#) describes the output fields for the **show mpls cspf** command. Output fields are listed in the approximate order in which they appear.

Table 97: show mpls cspf Output Fields

Field Name	Field Description
Queue length	Number of LSPs queued for automatic path computation.
current	Current queue length.
maximum	Maximum queue length (high-water mark).
dequeued	Number of aborted computation attempts.
Paths	Counters for label-switched path computations.
total	Sum of the next four fields.
successful	Number of path computations that were successfully completed.
no route	Number of path computations that failed because the destination is unreachable.
Sys Error	Number of path computations that failed because of lack of memory.
CSPFs	Total number of CSPF computations. A single path might require multiple CSPF computations.
Time	Time, in seconds, required to perform the label-switched path computation.
Total	Total amount of time consumed by the CSPF path computation algorithm.
CSPFs	Total number of CSPF computations.
Avg per CSPF	Average amount of time required for each CSPF computation.
% of rpd	Percentage of routing process CPU used in the CSPF computation.

Sample Output

show mpls cspf

user@host> **show mpls cspf**

CSPF statistics					
Queue length	current	maximum	dequeued		
	0	0	0		
Paths	total	successful	no route	sys error	CSPFs
	0	0	0	0	0
Time (secs)	total	CSPFs	avg per CSPF	% of rpd	
	0.000000	0.000000	0.000000	0.0000	

show mpls diffserv-te

List of Syntax

[Syntax on page 3197](#)

[Syntax \(EX Series Switches\) on page 3197](#)

Syntax

```
show mpls diffserve-te
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls diffserve-te
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.

Options

none—Display DiffServ classes and priorities used by MPLS LSPs.

instance *instance-name*—(Optional) Display DiffServ classes and priorities used by MPLS LSPs for the specified instance. If ***instance-name*** is omitted, DiffServ information 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.

Required Privilege Level

view

List of Sample Output

[show mpls diffserv-te on page 3198](#)

Output Fields

[Table 98 on page 3198](#) describes the output fields for the **show mpls diffserv-te** command. Output fields are listed in the approximate order in which they appear.

Field Name	Field Description
Bandwidth model	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
TE class	DiffServ traffic engineering class.
Traffic class	<p>MPLS class type that corresponds to the DiffServ traffic engineering class:</p> <ul style="list-style-type: none"> • ct0—Best effort • ct1—Assured forwarding • ct2—Expedited forwarding • ct3—Network control
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.

TE class	Traffic class	Priority
TE-LSR	TE-LSR	1
TE-LSR	TE-LSR	2
TE-LSR	TE-LSR	3
TE-LSR	TE-LSR	4
TE-LSR	TE-LSR	5
TE-LSR	TE-LSR	6
TE-LSR	TE-LSR	7
TE-LSR	TE-LSR	8
TE-LSR	TE-LSR	9
TE-LSR	TE-LSR	10
TE-LSR	TE-LSR	11
TE-LSR	TE-LSR	12
TE-LSR	TE-LSR	13
TE-LSR	TE-LSR	14
TE-LSR	TE-LSR	15
TE-LSR	TE-LSR	16
TE-LSR	TE-LSR	17
TE-LSR	TE-LSR	18
TE-LSR	TE-LSR	19
TE-LSR	TE-LSR	20
TE-LSR	TE-LSR	21
TE-LSR	TE-LSR	22
TE-LSR	TE-LSR	23
TE-LSR	TE-LSR	24
TE-LSR	TE-LSR	25
TE-LSR	TE-LSR	26
TE-LSR	TE-LSR	27
TE-LSR	TE-LSR	28
TE-LSR	TE-LSR	29
TE-LSR	TE-LSR	30
TE-LSR	TE-LSR	31
TE-LSR	TE-LSR	32
TE-LSR	TE-LSR	33
TE-LSR	TE-LSR	34
TE-LSR	TE-LSR	35
TE-LSR	TE-LSR	36
TE-LSR	TE-LSR	37
TE-LSR	TE-LSR	38
TE-LSR	TE-LSR	39
TE-LSR	TE-LSR	40
TE-LSR	TE-LSR	41
TE-LSR	TE-LSR	42
TE-LSR	TE-LSR	43
TE-LSR	TE-LSR	44
TE-LSR	TE-LSR	45
TE-LSR	TE-LSR	46
TE-LSR	TE-LSR	47
TE-LSR	TE-LSR	48
TE-LSR	TE-LSR	49
TE-LSR	TE-LSR	50
TE-LSR	TE-LSR	51
TE-LSR	TE-LSR	52
TE-LSR	TE-LSR	53
TE-LSR	TE-LSR	54
TE-LSR	TE-LSR	55
TE-LSR	TE-LSR	56
TE-LSR	TE-LSR	57
TE-LSR	TE-LSR	58
TE-LSR	TE-LSR	59
TE-LSR	TE-LSR	60
TE-LSR	TE-LSR	61
TE-LSR	TE-LSR	62
TE-LSR	TE-LSR	63
TE-LSR	TE-LSR	64
TE-LSR	TE-LSR	65
TE-LSR	TE-LSR	66
TE-LSR	TE-LSR	67
TE-LSR	TE-LSR	68
TE-LSR	TE-LSR	69
TE-LSR	TE-LSR	70
TE-LSR	TE-LSR	71
TE-LSR	TE-LSR	72
TE-LSR	TE-LSR	73
TE-LSR	TE-LSR	74
TE-LSR	TE-LSR	75
TE-LSR	TE-LSR	76
TE-LSR	TE-LSR	77
TE-LSR	TE-LSR	78
TE-LSR	TE-LSR	79
TE-LSR	TE-LSR	80
TE-LSR	TE-LSR	81
TE-LSR	TE-LSR	82
TE-LSR	TE-LSR	83
TE-LSR	TE-LSR	84
TE-LSR	TE-LSR	85
TE-LSR	TE-LSR	86
TE-LSR	TE-LSR	87
TE-LSR	TE-LSR	88
TE-LSR	TE-LSR	89
TE-LSR	TE-LSR	90
TE-LSR	TE-LSR	91
TE-LSR	TE-LSR	92
TE-LSR	TE-LSR	93
TE-LSR	TE-LSR	94
TE-LSR	TE-LSR	95
TE-LSR	TE-LSR	96
TE-LSR	TE-LSR	97
TE-LSR	TE-LSR	98
TE-LSR	TE-LSR	99
TE-LSR	TE-LSR	100

```
te0      ct0      3
```

tel	ct1	2
-----	-----	---

show mpls interface

Syntax

```
show mpls interface
```

Release Information

Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display information about MPLS-enabled interfaces. MPLS is enabled on an interface when the interface is configured with both the **set protocols mpls interface *interface-name*** and **set interfaces *interface-name* unit 0 family mpls** commands.

Required Privilege Level

view

RELATED DOCUMENTATION

- [Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) | 1318](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) | 1320](#)
- [Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

List of Sample Output

[show mpls interface on page 3200](#)

Output Fields

[Table 99 on page 3199](#) describes the output fields for the **show mpls interface** command. Output fields are listed in the approximate order in which they appear.

Table 99: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: Up or Dn (down).
Administrative groups	Administratively assigned colors of the link.

Sample Output

show mpls interface

user@switch> **show mpls interface**

Interface	State	Administrative groups
so-1/0/0.0	Up	Blue Yellow Red

show mpls egress-protection

Syntax

```
show mpls egress-protection  
<brief | detail>  
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 11.4R3.

Description

Display information about egress protection.

NOTE: Use this command on the device configured as the protector PE router to display information about egress protection. If you use this command on the device configured as the primary PE router, no output is displayed.

Options

none—Display standard information about egress protection.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

Example: Configuring MPLS Egress Protection for Layer 3 VPN Services

Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

List of Sample Output

[show mpls egress-protection detail \(Centralized Protector\) on page 3202](#)

[show mpls egress-protection detail \(Collocated Protector\) on page 3202](#)

Output Fields

Table 95 on page 3189 describes the output fields for the **show mpls egress-protection detail** command. Output fields are listed in the approximate order in which they appear.

Table 100: show mpls lsp Output Fields

Field Name	Field Description
Instance	Indicates egress instance name
Type	Indicates type of the VRF. It can be either local-vrf or remote-vrf
RIB	Indicates the edge-protection created routing table
Context-Id	Indicates the context-ID associated with the RIB.
Interface/Enhanced-lookup	<p>Show VT interfaces associated with the backup RIB.</p> <p>Shows Enhanced-lookup for MX Series 5G Universal Routing Platforms with the Enhanced IP Network Services mode configured using the network-services enhanced-ip statement at the [edit chassis] hierarchy level.</p>

Sample Output

show mpls egress-protection detail (Centralized Protector)

```
user@host> show mpls egress-protection detail
```

Instance	Type	Protection-Type
rsite1	remote-vrf	Protector
RIB __99.99.1.4-rsite1__.inet.0, Context-Id 99.99.1.4, Enhanced-lookup		
Route Target 1:1		
rsite24	remote-vrf	Protector
RIB __99.99.1.4-rsite24__.inet.0, Context-Id 99.99.1.4, Enhanced-lookup		
Route Target 100:1023		

Sample Output

show mpls egress-protection detail (Collocated Protector)

```
user@host> show mpls egress-protection detail
```



```
Instance          Type      Protection-Type
site2             local-vrf  Protector
  RIB __66.6.6.6-site2__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031809

  Route Target 100:251
site12            local-vrf  Protector
  RIB __66.6.6.6-site12__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031808

  Route Target 100:250
  Route Target 100:251

site2             local-vrf  Protector
  RIB __66.6.6.6-site2__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031809

  Route Target 100:251
```


show mpls interface

List of Syntax

[Syntax on page 3204](#)

[Syntax \(EX Series Switches\) on page 3204](#)

Syntax

```
show mpls interface
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show mpls interface
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.

Options

none—Display information about MPLS-enabled interfaces.

instance *instance-name*—(Optional) Display information about MPLS-enabled interfaces for the specified routing instance. If ***instance-name*** is omitted, information about MPLS-enabled interfaces 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.

Additional Information

MPLS is enabled on an interface when the interface is configured with both the **set protocol mpls interface *interface-name*** and **set interface *interface-name* unit 0 family mpls** statements.

Required Privilege Level

view

List of Sample Output

[show mpls interface on page 3206](#)

Output Fields

Table 101 on page 3205 describes the output fields for the **show mpls interface** command. Output fields are listed in the approximate order in which they appear.

Table 101: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: Up or Dn (down).
Administrative groups	Administratively assigned colors of the link.
Maximum labels	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the maximum-labels statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels.
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the protection-revert-time statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels.
Always mark connection protection tlv	Enabled or Disabled : Enabled indicates that the always-mark-connection-protection-tlv statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the switch-away-lsps statement must be configured.

Table 101: show mpls interface Output Fields (*continued*)

Field Name	Field Description
Switch away lsps	Enabled or Disabled : Enabled indicates that the switch-away-lsps statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

Sample Output

show mpls interface

user@host> show mpls interface

```
Interface: ge-0/2/1.57
  State: Up
  Administrative group: <none>
  Maximum labels: 5
  Static protection revert time: 5 seconds
  Always mark connection protection tlv: Disabled
  Switch away lsps : Disabled
```


show mpls label usage

Syntax

```
show mpls label usage
<label label value>
<label-range range-start range-end>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 15.1.

Support for the **label** statement added in Junos OS Release 17.2.

Support for the **label-range** statement added in Junos OS Release 17.2.

Description

Show the available label space resource in RPD and also the applications that use the label space in RPD. There are four different label spaces currently used in MPLS—namely LSI, dynamic, block, and static. Each label space has a fixed number and cannot grow beyond the fixed value. Using this command, the administrator can monitor the available labels in each label space and the applications that are using the labels. Based on the availability of labels, the administrator can decide to stop any service and free some labels or use other service where the labels are available.

Starting in Junos OS Release 17.2, you can configure the **enhanced-ip** command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can also separate the MPLS labels used for different label spaces which provides more flexibility and scalability.

When you set each member router's network services to **enhanced-ip**, only MPC or Modular Interface Cards (MICs) modules and Multiservices Dense Port Concentrator (MS-DPC) modules are powered on in the chassis. Non-service DPCs do not work with enhanced IP network services.

Options

none— Display the available labels in each label space and the applications using the labels.

label *label value*—(Optional) Display the information about which *label value* is used by which protocol, if any.

label-range *range-start range-end* —(Optional) Display the complete information about the **label-range** specified. With the **enhanced-ip** command enabled on the supported device, effective ranges and configured ranges along with details of different label spaces such as LSI, dynamic, block, and static types are displayed.

logical-system (*all* | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

Once the label space crosses the threshold, a new syslog message is added.

<label-space-name> label space usage crossed threshold limit of 90%.

For instance, LSI label space usage crossed threshold limit of 90%.

Required Privilege Level

view

List of Sample Output

[show mpls label usage on page 3208](#)

Output Fields

[Table 102 on page 3208](#) describes the output fields for the **show mpls label usage** command. Output fields are listed in the order in which they appear.

Table 102: show mpls label usage Fields

Field Name	Field Description
Label Space	Indicates the different types of labels currently used in MPLS.
Total	Indicates the total label space available.
Available	Indicates the number of freely available labels and also the percentage of the label space available.
Applications	Indicates the applications that use the MPLS label spaces.
Effective Ranges	Indicates actual ranges in use, which can be different from configured ranges, if conflicting with label already allocated.
Configured Ranges	Indicates the currently configured range assigned to different label spaces on the device.

Sample Output

show mpls label usage

user@host> **show mpls label usage**

Label space	Total	Available	Applications
LSI	999984	999971 (100.00%)	BGP/LDP VPLS with no-tunnel-services, BGP
L3VPN with vrf-table-label			


```

Block          999984  999971 (100.00%) BGP/LDP VPLS with tunnel-services, BGP L2VPN
Dynamic        999984  999971 (100.00%) RSVP, LDP, PW, L3VPN, RSVP-P2MP, LDP-P2MP,
MVPN, EVPN, BGP
Static         48576   48576  (100.00%) Static LSP, Static PW

```

With **enhanced-ip** enabled on the supported device, you get the following additional output.

user@host> **show mpls label usage**

```

Label space Total    Available      Applications
LSI          996983  996983 (100.00%) BGP/LDP VPLS with no-tunnel-services, BGP
L3VPN with vrf-table-label
Block        996983  996983 (100.00%) BGP/LDP VPLS with tunnel-services, BGP L2VPN
Dynamic      996983  996983 (100.00%) RSVP, LDP, PW, L3VPN, RSVP-P2MP, LDP-P2MP,
MVPN, EVPN, BGP
Static       48576   48576  (100.00%) Static LSP, Static PW
Effective Ranges
Range name  Shared with Start    End
Dynamic     16      999
Dynamic     4001    999999
Static      1000000 1048575
SRGB        1000    2999    OSPF
SRGB        3000    4000    GLOBAL
Configured Ranges
Range name  Shared with Start    End
Dynamic     16      999
Dynamic     4001    999999
Static      1000000 1048575
SRGB        1000    2999    OSPF
SRGB        3000    4000    GLOBAL

```

user@host> **show mpls label usage label 101**

```

Label 101 is used by protocol BGP

```

user@host> **show mpls label usage label 102**

```

Label 102 is used by protocol LDP

```

user@host> **show mpls label usage label 103**

Label 103 is not allocated to any protocol

show mpls label usage label-range

Syntax

```
show mpls label usage label-range
<block-label-range range-start range-end>
<dynamic-label-range range-start range-end>
<label-limit label-limit value>
<lsi-label-range range-start range-end>
<static-label-range range-start range-end>
```

Release Information

Command introduced in Junos OS Release 17.2.

Description

There are four different label spaces currently used in MPLS—namely LSI, dynamic, block, and static. Each label space has a fixed number and cannot grow beyond the fixed value. Using **show mpls label usage label-range** command, the administrator can monitor the available labels in each label space and the applications that are using the labels. Based on the availability of labels, the administrator can decide to stop any service and free some labels or use other service where the labels are available.

Starting in Junos OS Release 17.2, you can configure the **enhanced-ip** command, which is supported on platforms using Modular Port Concentrators (MPCs) equipped with Junos Trio chipsets. You can also separate the MPLS labels used for different label spaces which provides more flexibility and scalability.

When you set each member router's network services to **enhanced-ip**, only MPC or Modular Interface Cards (MICs) modules and Multiservices Dense Port Concentrator (MS-DPC) modules are powered on in the chassis. Non-service DPCs do not work with enhanced IP network services.

Options

block-label-range range-start range-end—Display the details of block label type.

dynamic-label-range range-start range-end —Display the details of dynamic label type.

label-limit label-limit value—Limit for the number of concurrent active labels.

lsi-label-range range-start range-end—Display the details of LSI label type.

static-label-range range-start range-end—Display the details of static label type.

Required Privilege Level

view

RELATED DOCUMENTATION

[show mpls label usage](#) | 3207

Output Fields

[Table 103 on page 3212](#) describes the output fields for the **show mpls label usage label-range** command. Output fields are listed in the order in which they appear.

Table 103: show mpls label usage label-range Fields

Field Name	Field Description
Label Space	Indicates the different types of labels currently used in MPLS.
Total	Indicates the total label space available.
Available	Indicates the number of freely available labels and also the percentage of the label space available.
Applications	Indicates the applications that use the MPLS label spaces.
Effective Ranges	Indicates actual ranges in use, which can be different from configured ranges, if conflicting with label already allocated.
Configured Ranges	Indicates the currently configured range assigned to different label spaces on the device.
Total	Indicates the currently used labels with label type and application type details.

Sample Output

With the **enhanced-ip** command enabled on the supported device, you get the following output.

```
user@host> show mpls label usage label-range 16 600
```

```

Label space Total   Available      Applications
LSI           101      99   (98.02% ) BGP/LDP VPLS with no-tunnel-services, BGP
L3VPN with vrf-table-label
Block        101      101  (100.00%) BGP/LDP VPLS with tunnel-services, BGP L2VPN
Dynamic       101      98   (97.03% ) RSVP, LDP, PW, L3VPN, RSVP-P2MP, LDP-P2MP,
MVPN, EVPN, BGP
Static       48576  48576  (100.00%) Static LSP, Static PW
Effective Ranges

```


Range name	Shared	with	Start	End
LSI	300		400	
Block	500		600	
Dynamic	100		200	
Static	1000000		1048575	

Configured Ranges

Range name	Shared	with	Start	End
LSI	300		400	
Block	500		600	
Dynamic	100		200	
Static	1000000		1048575	

Total(16 to 600) 5

Label type	Alloc	count
LSI	2	
Dynamic	3	

App type	Alloc	count
LDP	1	
BGP	2	
RT_INSTANCE	2	

show mpls lsp

List of Syntax

[Syntax on page 3214](#)

[Syntax \(EX Series Switches\) on page 3214](#)

Syntax

```
show mpls lsp
<brief | detail | extensive | terse>
<abstract-computation>
<autobandwidth>
<bidirectional | unidirectional>
<bypass>
<count-active-routes>
<defaults>
<descriptions>
<down | up>
<externally-controlled>
<externally-provisioned>
<instance routing-instance-name>
<locally-provisioned>
<logical-system (all | logical-system-name)>
<lsp-type>
<name name>
<p2mp>
<reverse-statistics>
<segment>
<statistics>
<transit>
```

Syntax (EX Series Switches)

```
show mpls lsp
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<descriptions>
<down | up>
<externally-controlled>
<externally-provisioned>
<lsp-type>
<name name>
<p2mp>
```



```
<statistics>
<transit>
```

Release Information

Command introduced before Junos OS Release 7.4.

defaults option added in Junos OS Release 8.5.

Command introduced in Junos OS Release 9.5 for EX Series switches.

autobandwidth option added in Junos OS Release 11.4.

externally-controlled option added in Junos OS Release 12.3.

externally-provisioned option added in Junos OS Release 13.3.

Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all configured and active dynamic MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

In the extensive command output, the duplicate back-to-back messages are recorded as aggregated messages. An additional timestamp is included for these aggregated messages, where if the aggregated messages are five or less, timestamp deltas are recorded for each message, and if the aggregated messages are greater than five, the first and last timestamp is recorded.

For example:

- All timestamps

```
9204 Jun 29 13:23:45.405 54.239.43.110: Explicit Route: bad strict route [3
times - 13:21:00, 13:22:01, 13:23:10]
```

- Timestamp deltas

```
9204 Jun 29 13:23:45.405 54.239.43.110: Explicit Route: bad strict route [3
times - 13:21:00, +1:01, +2:10]
```

- First and last timestamp


```
9204 Jun 29 13:23:45.405 54.239.43.110: Explicit Route: bad strict route [6
times - 13:21:00, 13:23:10]
```

abstract-computation—(Optional) Display abstract computation preprocessing for LSPs.

See [show mpls lsp abstract-computation](#) for more details.

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.

locally-provisioned—(Optional) Display LSPs that have been provisioned locally by the Path Computation Client (PCC).

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.

- **pop-and-forward**—Sessions that originate from RSVP-TE pop-and-forward LSP tunnels.
- **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.

reverse-statistics—(Optional) Display packet statistics for reverse direction of LSPs.

segment—(Optional) Display segment identifier (SID) labels.

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 | 3103](#)

[show mpls lsp autobandwidth | 3243](#)

List of Sample Output

[show mpls lsp defaults on page 3228](#)

[show mpls lsp descriptions on page 3228](#)

[show mpls lsp detail on page 3228](#)

[show mpls lsp detail \(When Egress Protection Is in Standby Mode\) on page 3229](#)

[show mpls lsp detail \(When Egress Protection Is in Effect During a Local Repair\) on page 3230](#)

[show mpls lsp extensive on page 3231](#)

[show mpls lsp ingress extensive on page 3233](#)

[show mpls lsp extensive \(automatic bandwidth adjustment enabled\) on page 3234](#)

[show mpls lsp bypass extensive on page 3235](#)

[show mpls lsp p2mp on page 3236](#)

[show mpls lsp p2mp detail on page 3236](#)

[show mpls lsp detail count-active-routes on page 3237](#)

[show mpls lsp statistics extensive on page 3238](#)

Output Fields

Table 104 on page 3218 describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

Table 104: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSP	Information about LSPs on the ingress routing device. Each session has one line of output.	All levels
Egress LSP	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
Transit LSP	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
P2MP name	Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format <i>identifier:vppls:router-id:routing-instance-name</i> . The <i>identifier</i> is automatically generated by Junos OS.	All levels
P2MP branch count	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
P	An asterisk (*) under this heading indicates that the LSP is a primary path.	All levels
address	(detail and extensive) Destination (egress routing device) of the LSP.	detail extensive
To	Destination (egress routing device) of the session.	brief
From	Source (ingress routing device) of the session.	brief detail
State	State of the LSP handled by this RSVP session: Up , Dn (down), or Restart .	brief detail

Table 104: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active Route	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail extensive
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	brief
ActivePath	(Ingress LSP) Name of the active path: Primary or Secondary .	detail extensive
LSPname	Name of the LSP.	brief detail
Statistics	Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation).	extensive
Aggregate statistics	Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the clear mpls lsp statistics command.	extensive
Packets	Displays the number of packets transmitted over the LSP.	brief extensive
Bytes	Displays the number of bytes transmitted over the LSP.	brief extensive
DiffServeInfo	Type of LSP: multiclass LSP (multiclass diffServ-TE LSP) or Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP).	detail
LSPtype	Type of LSP: <ul style="list-style-type: none"> • Static configured—Static • Dynamic configured—Dynamic • Externally controlled—External path computing entity Also indicates if the LSP is a Penultimate hop popping LSP or an Ultimate hop popping LSP.	detail extensive

Table 104: show mpls lsp Output Fields (continued)

Field Name	Field Description	Level of Output
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices.	detail
Bidir	(GMPLS) The LSP allows data to travel in both directions between GMPLS devices.	All levels
Bidirectional	(GMPLS) The LSP allows data to travel both ways between GMPLS devices.	All levels
FastReroute desired	Fast reroute has been requested by the ingress routing device.	detail
Link protection desired	Link protection has been requested by the ingress routing device.	detail
Node/Link protection desired	Link protection has been requested by the ingress routing device.	detail
LSP Control Status	<p>(Ingress LSP) LSP control mode:</p> <ul style="list-style-type: none"> • External—By default, all PCE-controlled LSPs are under external control. When an LSP is under external control, the PCC uses the PCE-provided parameters to set up the LSP. • Local—A PCE-controlled LSP can come under local control. When the LSP switches from external control to local control, path computation is done using the CLI-configured parameters and constraint-based routing. Such a switchover happens only when there is a trigger to re-signal the LSP. Until then, the PCC uses the PCE-provided parameters to signal the PCE-controlled LSP, although the LSP remains under local control. <p>A PCE-controlled LSP switches to local control from its default external control mode in cases such as no connectivity to a PCE or when a PCE returns delegation of LSPs back to the PCC.</p>	extensive

Table 104: show mpls lsp Output Fields (continued)

Field Name	Field Description	Level of Output
External Path CSPF status	(PCE-controlled LSPs) Status of the PCE-controlled LSP with per path attributes: <ul style="list-style-type: none"> • Local • External 	extensive
Externally Computed ERO	(PCE-controlled LSPs) Externally computed explicit route when the route object is not null or empty. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	extensive
EXTCTRL_LSP	(PCE-controlled LSPs) Display path history including the bandwidth, priority, and metric values received from the external controller.	extensive
flap counter	Counts the number of times a LSP flaps down or up.	extensive
LoadBalance	(Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: Most-fill , Least-fill , or Random .	detail extensive
Signal type	Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: DS0 , DS3 , STS-1 , STM-1 , or STM-4 .	All levels
Encoding type	LSP encoding type: Packet , Ethernet , PDH , SDH/SONET , Lambda , or Fiber .	All levels
Switching type	Type of switching on the links needed for the LSP: Fiber , Lambda , Packet , TDM , or PSC-1 .	All levels
GPID	Generalized Payload Identifier (identifier of the payload carried by an LSP): HDLC , Ethernet , IPv4 , PPP , or Unknown .	All levels
Protection	Configured protection capability desired for the LSP: Extra , Enhanced , none , One plus one , One to one , or Shared .	All levels
Upstream label in	(Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.	All levels
Upstream label out	(Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.	All levels
Suggested label received	(Bidirectional LSPs) Label the upstream interface suggests to use in the Resv message that is sent.	All levels

Table 104: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Suggested label sent	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
Autobandwidth	(Ingress LSP) The LSP is performing autobandwidth allocation.	detail extensive
Mbb counter	Counts the number of times a LSP incurs MBB.	extensive
MinBW	(Ingress LSP) Configured minimum value of the LSP, in bps.	detail extensive
MaxBW	(Ingress LSP) Configured maximum value of the LSP, in bps.	detail extensive
Dynamic MinBW	(Ingress LSP) Displays the current dynamically specified minimum bandwidth allocation for the LSP, in bps.	detail extensive
Dynamic MinBW	(Ingress LSP) Displays the current dynamically specified minimum bandwidth allocation for the LSP, in bps.	detail extensive
AdjustTimer	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	detail extensive
Adjustment Threshold	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	detail extensive
Time for Next Adjustment	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	detail extensive
Time of Last Adjustment	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	detail extensive
MaxAvgBW util	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	detail extensive
Overflow limit	(Ingress LSP) Configured value of the threshold overflow limit.	detail extensive
Overflow sample count	(Ingress LSP) Current value for the overflow sample count.	detail extensive

Table 104: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Bandwidth Adjustment in <i>nnn</i> second(s)	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	detail extensive
Underflow limit	(Ingress LSP) Configured value of the threshold underflow limit.	detail extensive
Underflow sample count	(Ingress LSP) Current value for the underflow sample count.	detail extensive
Underflow Max AvgBW	(Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.	detail extensive
Active path indicator	(Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path. *Primary long Standby short	detail extensive
Primary	(Ingress LSP) Name of the primary path.	detail extensive
Secondary	(Ingress LSP) Name of the secondary path.	detail extensive
Standby	(Ingress LSP) Name of the path in standby mode.	detail extensive
State	(Ingress LSP) State of the path: Up or Dn (down).	detail extensive
COS	(Ingress LSP) Class-of-service value.	detail extensive
Bandwidth per class	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	detail extensive
Priorities	(Ingress LSP) Configured value of the setup priority and the hold priority respectively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority).	detail extensive
OptimizeTimer	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive

Table 104: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
SmartOptimizeTimer	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
Reoptimization in xxx seconds	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	detail extensive
Computed ERO (S [L] denotes strict [loose] hops)	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	detail extensive
CSPF metric	(Ingress LSP) Constrained Shortest Path First metric for this path.	detail extensive

Table 104: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Received RRO	<p>(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If Received RRO is different from Computed ERO, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> • 0x01—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding Path message. • 0x02—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • 0x03—Combination of 0x01 and 0x02. • 0x04—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • 0x08—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the Local protection available bit is set but the Node protection bit is cleared. • 0x09—Detour is established. Combination of 0x01 and 0x08. • 0x10—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • 0x20—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently. • 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08. 	detail extensive
Labels	<p>Labels of pop-and-forward LSP tunnel:</p> <ul style="list-style-type: none"> • P—Pop labels. • D—Delegation labels. 	extensive
Index number	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	extensive

Table 104: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Date	(Ingress LSP) Date of the LSP event.	extensive
Time	(Ingress LSP) Time of the LSP event.	extensive
Event	(Ingress LSP) Description of the LSP event.	extensive
Created	(Ingress LSP) Date and time the LSP was created.	extensive
Resv style	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail extensive
Labelin	Incoming label for this LSP.	brief detail
Labelout	Outgoing label for this LSP.	brief detail
LSPname	Name of the LSP.	brief detail
Time left	Number of seconds remaining in the lifetime of the reservation.	detail
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
Port number	Protocol ID and sender or receiver port used in this RSVP session.	detail
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	detail
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	detail
RESV rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the PATH rcvfrom field, indicates that the RSVP negotiation is complete.	detail

Table 104: show mpls lsp Output Fields (continued)

Field Name	Field Description	Level of Output
Record route	Recorded route for the session, taken from the record route object.	detail
Pop-and-forward	Attributes of the pop-and-forward LSP tunnel.	extensive
ETLD In	Number of transport labels that the LSP-Hop can potentially receive from its upstream hop. It is recorded as Effective Transport Label Depth (ETLD) at the transit and egress devices.	extensive
ETLD Out	Number of transport labels the LSP-Hop can potentially send to its downstream hop. It is recorded as ETLD at the transit and ingress devices.	extensive
Delegation hop	Specifies if the transit hop is selected as a delegation label: <ul style="list-style-type: none"> • Yes • No 	extensive
Soft preempt	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	detail
Soft preemption pending	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	detail
MPLS-TE LSP Defaults	Default settings for MPLS traffic engineered LSPs: <ul style="list-style-type: none"> • LSP Holding Priority—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. • LSP Setup Priority—Determines whether a new LSP that preempts an existing LSP can be established. • Hop Limit—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress). • Bandwidth—Specifies the bandwidth in bits per second for the LSP. • LSP Retry Timer—Length of time in seconds that the ingress router waits between attempts to establish the primary path. 	defaults

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth**. You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.

Sample Output

show mpls lsp defaults

user@host> show mpls lsp defaults

```
MPLS-TE LSP Defaults
  LSP Holding Priority      0
  LSP Setup Priority       7
  Hop Limit                255
  Bandwidth                0
  LSP Retry Timer          30 seconds
```

show mpls lsp descriptions

user@host> show mpls lsp descriptions

```
Ingress LSP: 3 sessions
To           LSP name           Description
10.0.0.195   to-sanjose                 to-sanjose-desc
10.0.0.195   to-sanjose-other-desc      other-desc
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail

user@host> show mpls lsp detail

```
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPTYPE: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
  10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
      10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0
```



```

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp detail (When Egress Protection Is in Standby Mode)

user@host> show mpls lsp detail

```

Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, 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
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
Egress protection PLR as protector: In Use
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp extensive

```
user@host> show mpls lsp extensive
```

```

Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPTtype: Static Configured, Ultimate hop popping
  LSP Control Status: Externally controlled

```



```

LoadBalance: Random
Metric: 10
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary                               State: Up
  Priorities: 7 0
  External Path CSPF status: local
  Bandwidth: 98.76kbps
  SmartOptimizeTimer: 180
  Include All: green
  Externally Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 0)
1.2.3.2 S 2.3.3.2 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
  9 May 17 16:55:06.574 EXTCTRL LSP: Sent Path computation request and LSP status

  8 May 17 16:55:06.574 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 98760 req BW 0 admin group(exclude 0 include any 0 include all 16)
priority setup 5 hold 4 hops: 1.2.3.2 2.3.3.2
  7 May 17 16:55:06.574 Selected as active path
  6 May 17 16:55:06.558 EXTCTRL LSP: Sent Path computation request and LSP status

  8 May 17 16:55:06.574 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 98760 req BW 0 admin group(exclude 0 include any 0 include all 16)
priority setup 5 hold 4 hops: 1.2.3.2 2.3.3.2
  7 May 17 16:55:06.574 Selected as active path
  6 May 17 16:55:06.558 EXTCTRL LSP: Sent Path computation request and LSP status

  5 May 17 16:55:06.558 EXTCTRL_LSP: Computation request/lsp status contains:
signalled bw 98760 req BW 0 admin group(exclude 0 include any 0 include all 16)
priority setup 5 hold 4 hops: 1.2.3.2 2.3.3.2
  4 May 17 16:55:06.557 Record Route: 1.2.3.2 2.3.3.2
  3 May 17 16:55:06.557 Up
  2 May 17 16:55:06.382 Originate Call
  1 May 17 16:55:06.382 EXTCTRL_LSP: Received setup parameters :: local_cspf,
1.2.3.2 2.3.3.2
  Created: Tue May 17 16:55:07 2016
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>

```

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)
  LSPTtype: Static Pop-and-forward Configured, Penultimate hop popping
  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):
        (Labels: P=Pop D=Delegation)
        80.1.1.2(Label=18 P) 50.1.1.2(Label=17 P) 70.1.1.2(Label=16 P)
        92.1.1.1(Label=16 D) 93.1.1.2(Label=16 P) 99.1.1.1(Label=16 P)
        99.2.1.1(Label=16 P) 99.3.1.2(Label=3)
    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
LSPTtype: 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
Explot 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)
  LSPTtype: 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 abstract-computation

Syntax

```
show mpls lsp abstract-computation
<brief | detail | extensive>;
<logical-system (all | logical-system-name)>;
<name lsp-name>;
```

Release Information

Command introduced in Junos OS Release 17.1 for all platforms.

Description

Display the ingress to egress abstract hop computation used by the constrained shortest path in the preprocessing for LSPs. The command output displays the various computation passes involved per LSP, and the qualifying exit devices for each pass. It also displays the affinity per pass, and the current start device chosen for the pass.

Options

brief | detail | extensive—(Optional) Display the desired level of output.

logical-system (all | *logical-system-name*)—(Optional) Display the abstract computation for abstract hop constraints on all logical systems or on a particular logical system.

lsp-name—(Optional) Name of the LSP for which the abstract hop computation is displayed.

Required Privilege Level

view

RELATED DOCUMENTATION

[Example: Configuring Abstract Hops for MPLS LSPs | 446](#)

[abstract-hop | 2534](#)

[constituent-list | 2578](#)

[show mpls abstract-hop-membership | 3170](#)

List of Sample Output

[show mpls lsp abstract-computation on page 3241](#)

Output Fields

[Table 105 on page 3241](#) describes the output fields for the **show mpls lsp abstract-computation** command. Output fields are listed in the approximate order in which they appear.

Table 105: show mpls lsp abstract-computation Output Fields

Field Name	Field Description
Path computation using abstract hops for LSP	Name of the LSP for which the abstract hop computation is performed.
Path type	The type of the path can be primary or secondary.
Path name	Name of the path.
Credibility	Credibility value associated with the interior gateway protocol in use.
Total no of CSPF passes	Number of constrained shortest path passes for the abstract hop.
CSPF pass no	Constrained shortest path pass number for the abstract hop computation.
Start address of the pass	IP address where the pass starts.
Affinity	Name of the abstract hop.
Destination	Destination IP address for a node in the pass.
State	State of the backtracking: <ul style="list-style-type: none"> • Valid • Disqualified

Sample Output

show mpls lsp abstract-computation

user@R0> **show mpls lsp abstract-computation**

```

Path computation using abstract hops for LSP: R0-R31
Path type: Primary, Path name: prim
Credibility: 0, Total no of CSPF passes: 2
CSPF pass no: 0
Start address of the pass: 127.0.0.6
Destination: 127.0.0.1, State: VALID
Destination: 127.0.0.2, State: VALID
Destination: 127.0.0.3, State: VALID
Affinity: ahl

```



```
CSPF pass no: 1
Start address of the pass: 127.0.0.1
Destination: 127.0.0.3, State: VALID
Path type: Secondary, Path name: nonstdby
Path type: Standby, Path name: stdby
Credibility: 0, Total no of CSPF passes: 2
CSPF pass no: 0
Start address of the pass: 127.0.0.6
Destination: 127.0.0.3, State: VALID
Destination: 127.0.0.4, State: VALID
Affinity: ah2
CSPF pass no: 1
Start address of the pass: 127.0.0.4
Destination: 127.0.0.3, State: VALID
```


show mpls lsp autobandwidth

Syntax

```
show mpls lsp autobandwidth
<brief | detail | extensive>
<logical-system (all | logical-system-name)>
<name lsp-name>
```

Release Information

Statement introduced in Junos OS Release 11.4.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 15.1X54-D60 for the ACX5000 Series.

Statement introduced in Junos OS Release 17.2R1 for QFX10000 Series switches.

name *lsp-name* option introduced in Junos OS Release 18.1R1 for all platforms.

Description

Display automatic bandwidth information for the LSP(s).

After a Routing Engine switchover, the output of the **show mpls autobandwidth** command might not be up-to-date, as the automatic bandwidth information for the LSP(s) is gathered by the new master Routing Engine during the first adjustment interval.

Options

brief | detail | extensive —(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name *lsp-name*—(Optional) Specify name of the LSP for which the automatic bandwidth information should be displayed.

Required Privilege Level

view

RELATED DOCUMENTATION

[show mpls lsp | 3214](#)

[Achieving a Make-Before-Break, Hitless Switchover for LSPs | 538](#)

List of Sample Output

[show mpls lsp autobandwidth on page 3245](#)

Output Fields

Table 106 on page 3244 describes the output fields for the **show mpls lsp autobandwidth** command. Output fields are listed in the approximate order in which they appear.

Table 106: show mpls lsp autobandwidth Output Fields

Field Name	Field Description	Level of Output
To	Destination (egress routing device) of the session.	All Levels
From	Source (ingress routing device) of the session.	All Levels
LSPname	Name of the LSP.	All Levels
Min BW	(Ingress LSP) Configured minimum value of the LSP, in bps.	detail extensive
Max BW	(Ingress LSP) Configured maximum value of the LSP, in bps.	detail extensive
Max AvgBW util	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps. NOTE: In calculating this value, the sample collected during make before break (MBB) is ignored to prevent inaccurate results. The first sample after a bandwidth adjustment, or after a change in the LSP ID (regardless of path change), is also ignored.	detail extensive
Overflow limit	(Ingress LSP) Configured value of the threshold overflow limit.	detail extensive
Overflow sample count	(Ingress LSP) Current value for the overflow sample count.	detail extensive
Underflow limit	(Ingress LSP) Configured value of the threshold underflow limit.	detail extensive
Underflow sample count	(Ingress LSP) Current value for the underflow sample count.	detail extensive
Adjustment Timer	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	detail extensive
Adjustment Threshold	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	detail extensive
Time for Next Adjustment	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	detail extensive

Table 106: show mpls lsp autobandwidth Output Fields (*continued*)

Field Name	Field Description	Level of Output
Time of Last Adjustment	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	detail extensive
Last BW	Previous active bandwidth of the LSP.	detail extensive
Last Requested BW	Bandwidth requested in the previous automatic bandwidth adjustment.	detail extensive
Last Signaled BW	Bandwidth signaled in the previous automatic bandwidth adjustment.	detail extensive
Highest Watermark BW	Maximum bandwidth used by the LSP.	detail extensive
Total AutoBw Adjustments	Total number of attempts to adjust automatic bandwidth including failed and successful adjustments.	detail extensive
Successful Adjustments	Number of successful automatic bandwidth adjustments.	detail extensive
Failed Adjustments	Number of failed automatic bandwidth adjustments.	detail extensive

Sample Output

show mpls lsp autobandwidth

user@host> **show mpls lsp autobandwidth extensive**

```
To: 10.255.106.133,
From: 10.255.106.135, LSPname: r0-r1
Min BW: 100kbps, Max BW: 0bps, Max AvgBW util: 2.33249Mbps
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0
Adjustment Timer: 300 sec, Adjustment Threshold: 0
Time for Next Adjustment: 23 sec, Time of Last Adjustment: Fri Jun  3 21:05:37
2011
Last BW: 100kbps, Last Requested BW: 2.2169Mbps, Last Signaled BW: 2.2169Mbps,
Highest Watermark BW: 2.33249Mbps
Total AutoBw Adjustments: 1, Successful Adjustments: 1, Failed Adjustments: 0
```


show mpls path

List of Syntax

[Syntax on page 3246](#)

[Syntax \(EX Series Switches\) on page 3246](#)

Syntax

```
show mpls path
<instance instance-name>
<logical-system (all | logical-system-name)>
<path-name>
```

Syntax (EX Series Switches)

```
show mpls path
<path-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all MPLS LSPs.

instance *instance-name*—(Optional) Display the dynamic MPLS LSP for the specified instance. If *instance-name* is omitted, dynamic MPLS LSP for the master instance is displayed.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

path-name—(Optional) Display information about the specified LSP only.

Required Privilege Level

view

List of Sample Output

[show mpls path on page 3247](#)

Output Fields

Table 107 on page 3247 describes the output fields for the **show mpls path** command. Output fields are listed in the approximate order in which they appear.

Table 107: show mpls path Output Fields

Field Name	Field Description
Path name	Information about ingress LSPs. Each path has one line of output.
Address	Addresses of the routing devices that form the LSP.
Strict/loose address	Whether the address is configured as a strict or loose address.

Sample Output

show mpls path

user@host> **show mpls path**

Path name	Address	Strict/loose address
p1	123.456.55.6	Strict
	123.456.1.6	Loose
p2	191.456.1.4	Strict

show mpls srlg

Syntax

```
show mpls srlg
<logical-systems (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 11.4.

Description

Display Shared Risk Link Group (SRLG) cost and value configuration information.

NOTE: If an SRLG is associated with a link that is used by an ingress LSP in the router, then on deleting the SRLG configuration from that router, the SRLG gets removed from the SRLG table only on the next reoptimization of the LSP. Until then, the output of the **run show mpls srlg** command displays **Unknown-XXX** instead of the SRLG name and a non zero srlg-cost for that SRLG.

Options

logical-system (all | *logical-system-name*)—(Optional) View SRLG configuration information for all logical systems or a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[Example: Configuring SRLG | 278](#)

Output Fields

[Table 108 on page 3248](#) lists the output fields for the **show mpls srlg** command. Output fields are listed in the approximate order in which they appear.

Table 108: show mpls srlg Output Fields

Field Name	Field Description
SRLG	Name of the SRLG.

Table 108: show mpls srlg Output Fields (continued)

Field Name	Field Description
Value	A group ID for the SRLG ranging from 1 through 4294967295 .
Cost	A cost for the Shared Risk Link Group (SRLG) ranging from 1 through 65535 .

Sample Output

user@host> show mpls srlg

SRLG	Value	Cost
srlg-a	101	10

show mpls static-lsp

Syntax

```
show mpls static-lsp
<brief | detail | extensive | terse>
<bypass>
<descriptions>
<down | up>
<ingress>
<instance instance-name>
<logical-system (all | logical-system-name)>
<lsp-type>
<name name>
<statistics>
<transit>
```

Release Information

Command introduced in Junos OS Release 10.1.

Command updated in Junos OS Release 14.1X53-D25 to accommodate the stitching feature of MPLS.

Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Statement introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

Description

Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all configured and active static MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.

bypass—(Optional) Display LSPs used for protecting other static LSPs.

descriptions—(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the [edit protocols mpls static-label-switched-path *path-name* bypass], [edit protocols mpls static-label-switched-path *path-name* ingress], or [edit protocols mpls static-label-switched-path *path-name* transit *incoming-label*] hierarchy levels. Only static LSPs with a description are displayed.

down | up—(Optional) Display only static LSPs that are inactive or active, respectively.

instance *instance-name*—(Optional) Display information about all configured and active static MPLS LSPs for the specified routing instance. If ***instance-name*** is omitted, information about all configured and active static MPLS LSPs for the master instance is displayed.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified static LSP or group of LSPs.

statistics—(Optional) Display accounting information about static LSPs.

transit—(Optional) Display static LSPs transiting this routing device.

Required Privilege Level

view

List of Sample Output

[show mpls static-lsp extensive on page 3252](#)

[show mpls static-lsp statistics ingress on page 3253](#)

[show mpls static-lsp \(when MPLS stitching is used\) on page 3253](#)

Output Fields

[Table 94 on page 3180](#) describes the output fields for the **show mpls static-lsp** command. Output fields are listed in the approximate order in which they appear.

Table 109: show mpls static-lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSPs	Information about the static LSPs on the ingress routing device. Each session has one line of output.	All levels
Transit LSPs	Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
Bypass LSPs	Information about the bypass LSPs configured on the routing device. Each session has one line of output.	All levels
LSPname	Name of the static LSP.	All levels

Table 109: show mpls static-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
To	Destination (egress routing device) of the session.	All levels
State	State of the static LSP handled by this RSVP session: Up , Dn (down), or Restart .	All levels
Packets	Number of packet transiting the static LSP (statistics option only).	All levels
Bytes	Number of bytes transiting the static LSP (statistics option only).	All levels
Nexthop	IP address for the next-hop router for the static LSP.	detail, extensive
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
Link protection desired	Link protection has been requested by the ingress routing device.	detail, extensive
LabelOperation	Label operation to perform: Push , Pop , Swap .	detail, extensive
Outgoing-label	Outgoing label to use for the MPLS packet in either push or swap label operations.	detail, extensive
Created	(Ingress LSP) Date and time the static LSP was created.	extensive
Bandwidth	Bandwidth configured for the static LSP.	detail, extensive
Resv style	(Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	All levels

Sample Output

show mpls static-lsp extensive

user@host> **show mpls static-lsp extensive**

```
Ingress LSPs:
LSPname: alpha-to-beta, To: 192.168.14.1
State: Dn
```



```

Nexthop: 192.168.10.1
LabelOperation: Push, Outgoing-label: 1000001
Created: Thu Jan 14 16:44:43 2010
Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0

```

show mpls static-lsp statistics ingress

```
user@host> show mpls static-lsp statistics ingress
```

```

Ingress LSPs:
LSPname                To           State    Packets    Bytes
alpha-to-beta          192.168.14.1 Dn        NA         NA
Total 1, displayed 1, Up 0, Down 1

```

show mpls static-lsp (when MPLS stitching is used)

The show mpls static-lsp command was extended in Junos release 14.1X53-D25 to accommodate the stitching feature of MPLS. This example shows the LSP state as 'InProgress' because the LSP is waiting for protocol next-hop resolution. For more information, see

```
user@host> show mpls static-lsp
```

```

Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0

Transit LSPs: LSPname                Incoming-label  State
to-165                1000000        InProgress

```


show performance-monitoring mpls lsp

Syntax

```
show performance-monitoring mpls lsp  
<brief | detail | extensive>  
<name lsp name>
```

Release Information

Command introduced in Junos OS Release 15.1.

Description

Display the following performance monitoring data:

- Packet loss measurement
- Packet throughput measurement
- Two-way channel delay
- Round-trip delay
- Inter-packet delay variation (IPDV)

Options

none—Display standard information performance monitoring data.

brief | detail | extensive—(Optional) Display the specified level of output.

NOTE: The extensive option displays the same information as the detail option.

name *lsp name*—(Optional) Display information about the specified LSP.

Required Privilege Level

View

RELATED DOCUMENTATION

[clear performance-monitoring mpls lsp | 3107](#)

[performance-monitoring \(Protocols MPLS\) | 2727](#)

List of Sample Output

[show performance-monitoring mpls lsp on page 3258](#)

[show performance-monitoring mpls lsp detail on page 3259](#)

Output Fields

[Table 110 on page 3255](#) describes the output fields for the **show performance-monitoring mpls lsp** command.

Output fields are listed in the approximate order in which they appear.

Table 110: show performance-monitoring mpls lsp Output Fields

Field Name	Display Data	Field Description	Level of Output
Session	Total	Total number of performance monitoring sessions created.	All Levels
	Up	Number of performance monitoring sessions that are up and running.	All Levels
	Down	Number of performance monitoring sessions that are down.	All Levels
LSP name		Name of the LSP.	All Levels

Table 110: show performance-monitoring mpls lsp Output Fields (*continued*)

Field Name	Display Data		Field Description	Level of Output
Loss measurement Data	Traffic-class		Traffic class for which loss measurement is performed.	All Levels
	Queries sent		Total number of queries sent for loss measurement.	All Levels
	Responses received		Total number of responses received for loss measurement queries.	All Levels
	Responses dropped due to errors		Total number of loss measurement responses dropped due to errors.	All Levels
	Queries timeout		Number of timed out queries sent for loss measurement.	All Levels
	Forward loss measurement	Average packet loss	Average packet loss (total loss of packets divided by the total number of samples used since the session is up).	All Levels
		Average packet throughput	Total number of packets sent divided by the time considered for measurement.	All Levels
	Reverse loss measurement	Average packet loss	Average packet loss (total loss of packets divided by the total number of samples used since the session is up).	All Levels
		Average packet throughput	Total number of packets sent divided by the time considered for measurement.	All Levels

Table 110: show performance-monitoring mpls lsp Output Fields (*continued*)

Field Name	Display Data	Field Description	Level of Output
Delay measurement Data	Traffic-class	Traffic class for which delay measurement is performed.	All Levels
	Queries sent	Total number of queries sent for delay measurement.	All Levels
	Responses received	Total number of responses received for delay measurement queries.	All Levels
	Responses dropped due to errors	Total number of delay measurement responses dropped due to errors.	All Levels
	Queries timeout	Number of timed out queries sent for delay measurement.	All Levels
	Best 2-way channel delay	Best available two-way channel delay.	All Levels
	Worst 2-way channel delay	Worst available two-way channel delay.	All Levels
	Best round trip time	Best available round-trip time.	All Levels
	Worst round trip time	Worst available round-trip time.	All Levels
	Avg absolute fw delay variation	Average of the variation in forward delay.	All Levels
	Avg absolute rv delay variation	Average of the variation in reverse delay.	All Levels
	Two-way channel delay	Sum of packet delays, excluding the processing time of the remote provider edge (PE) router.	detail, extensive
	Two-way round trip delay	Total time taken for completing round-trip of packet.	detail, extensive

Sample Output

show performance-monitoring mpls lsp

user@host> **show performance-monitoring mpls lsp**

```

Session Total: 3 Up: 3 Down: 0
  LSP name:to_bad, PM State:Up
    Loss measurement Data:
      Duration: 00:04:43
      Traffic-class: None
      Queries sent: 282
      Responses received: 282
      Responses dropped due to errors: 0
      Queries timeout: 0
      Forward loss measurement:
        Average packet loss: 0
        Average packet throughput: 554338
      Reverse loss measurement:
        Average packet loss: 0
        Average packet throughput: 1352077
  LSP name:to_bad, PM State:Up
    Delay measurement Data:
      Duration: 00:04:43
      Traffic-class: 0
      Queries sent: 282
      Responses received: 282
      Responses dropped due to errors: 0
      Queries timeout: 0
      Best 2-way channel delay: 72 usecs
      Worst 2-way channel delay: 365 usecs
      Best round trip time: 843 usecs
      Worst round trip time: 105523 usecs
      Avg absolute fw delay variation: 1619 usecs
      Avg absolute rv delay variation: 1619 usecs
  LSP name:to_bad, PM State:Up
    Loss measurement Data:
      Duration: 00:04:43
      Traffic-class: None
      Queries sent: 282
      Responses received: 282
      Responses dropped due to errors: 0
      Queries timeout: 0
      Forward loss measurement:
        Average packet loss: 0

```



```

    Average packet throughput: 553927
Reverse loss measurement:
    Average packet loss: 0
    Average packet throughput: 1351531
Delay measurement Data:
    Best 2-way channel delay: 76 usecs
    Worst 2-way channel delay: 368 usecs
    Best round trip time: 1082 usecs
    Worst round trip time: 126146 usecs
    Avg absolute fw delay variation: 1618 usecs
    Avg absolute rv delay variation: 1619 usecs

```

show performance-monitoring mpls lsp detail

user@host> show performance-monitoring mpls lsp detail

```

Session Total: 3 Up: 3 Down: 0
LSP name:to_bad, PM State:Up
Loss measurement Data:
    Duration: 00:04:53
    Traffic-class: None
    Queries sent: 292
    Responses received: 292
    Responses dropped due to errors: 0
    Queries timeout: 0
Forward loss measurement:
    Average packet loss: 0
    Average packet throughput: 554486
    Packet loss samples:
        00000000 00000000 00000000 00000000 00000000
    Packet throughput samples:
        00554002 00557550 00557717 00558822 00557107
Reverse loss measurement:
    Average packet loss: 0
    Average packet throughput: 1352406
    Packet loss samples:
        00000000 00000000 00000000 00000000 00000000
    Packet throughput samples:
        01351088 01365948 01353926 01362976 01358788
LSP name:to_bad, PM State:Up
Delay measurement Data:
    Duration: 00:04:53
    Traffic-class: 0
    Queries sent: 292

```



```

Responses received: 292
Responses dropped due to errors: 0
Queries timeout: 0
Best 2-way channel delay: 72 usecs
Worst 2-way channel delay: 365 usecs
Best round trip time: 843 usecs
Worst round trip time: 105523 usecs
Avg absolute fw delay variation: 1683 usecs
Avg absolute rv delay variation: 1684 usecs
  Two-way channel delay:
    73 usecs 73 usecs 73 usecs 73 usecs 72 usecs
  Two-way round trip delay:
    922 usecs 2234 usecs 884 usecs 1121 usecs 1169 usecs
LSP name:to_bad, PM State:Up
Loss measurement Data:
  Duration: 00:04:53
  Traffic-class: None
  Queries sent: 292
  Responses received: 292
  Responses dropped due to errors: 0
  Queries timeout: 0
  Forward loss measurement:
    Average packet loss: 0
    Average packet throughput: 554089
    Packet loss samples:
      00000000 00000000 00000000 00000000 00000000
    Packet throughput samples:
      00554007 00557548 00557713 00558547 00557385
  Reverse loss measurement:
    Average packet loss: 0
    Average packet throughput: 1351914
    Packet loss samples:
      00000000 00000000 00000000 00000000 00000000
    Packet throughput samples:
      01358923 01352980 01362436 01223841 01496977
Delay measurement Data:
  Best 2-way channel delay: 76 usecs
  Worst 2-way channel delay: 368 usecs
  Best round trip time: 1082 usecs
  Worst round trip time: 126146 usecs
  Avg absolute fw delay variation: 1682 usecs
  Avg absolute rv delay variation: 1683 usecs
  Two-way channel delay:
    76 usecs 76 usecs 76 usecs 77 usecs 77 usecs

```


Two-way round trip delay:

107496 usecs 102369 usecs 104048 usecs 1433 usecs 103306 usecs

show route forwarding-table

Syntax

```
show route forwarding-table
<detail | extensive | summary>
<ccc ccc-interface-name>
<destination>
<family family-name>
<label label>
<matching ip_prefix>
<multicast>
<vpn vpn>
```

Release Information

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.

Command introduced in Junos OS Release 14.1X53-D30 for QFX Virtual Chassis and Virtual Chassis Fabric.

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.

Options

none—Display the routes in the forwarding table.

detail | extensive | summary—(Optional) Display the specified level of output.

ccc—(Optional) Display the specified circuit cross-connect interface name for entries to match.

destination—(Optional) Display the destination prefix.

family family-name—(Optional) Display routing table entries for the specified family: **ethernet-switching**, **inet**, **inet6**, **iso**, **mpls**, **vlan classification**.

label label—(Optional) Display route entries for the specified label name.

matching ip_prefix—(Optional) Display route entries for the specified IP prefix.

multicast—(Optional) Display route entries for multicast routes.

vpn vpn—(Optional) Display route entries for the specified VPN.

Required Privilege Level

view

RELATED DOCUMENTATION

Example: Configuring MPLS on EX8200 and EX4500 Switches 59
Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure) 99

List of Sample Output

- [show route forwarding-table on page 3266](#)
- [show route forwarding-table summary on page 3267](#)
- [show route forwarding-table extensive on page 3267](#)
- [show route forwarding-table ccc on page 3269](#)
- [show route forwarding-table family \(MPLS\) on page 3269](#)
- [show route forwarding-table family \(IPv6\) on page 3270](#)
- [show route forwarding-table label on page 3271](#)
- [show route forwarding-table matching on page 3271](#)
- [show route forwarding-table multicast on page 3271](#)

Output Fields

[Table 111 on page 3263](#) 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 111: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Routing table	Name of the routing table (for example, inet , inet6 , mpls).	All levels
Address family	Address family (for example, IP , IPv6 , ISO , MPLS).	All levels
Destination	Destination of the route.	detail , extensive

Table 111: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Route Type (Type)	<p>How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses):</p> <ul style="list-style-type: none"> • cloned (clon)—(TCP or multicast only) Cloned route. • destination (dest)—Remote addresses directly reachable through an interface. • destination down (iddn)—Destination route for which the interface is unreachable. • interface cloned (ifcl)—Cloned route for which the interface is unreachable. • route down (ifdn)—Interface route for which the interface is unreachable. • ignore (ignr)—Ignore this route. • interface (intf)—Installed as a result of configuring an interface. • permanent (perm)—Routes installed by the kernel when the routing table is initialized. • user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route reference (RtRef)	Number of routes to reference.	detail, extensive
Flags	<p>Route type flags:</p> <ul style="list-style-type: none"> • none—No flags are enabled. • accounting—Route has accounting enabled. • cached—Cache route. • incoming-iface interface-number —Check against incoming interface. • prefix load balance—Load balancing is enabled for this prefix. • sent to PFE—Route has been sent to the Packet Forwarding Engine. • static—Static route. 	extensive
Nexthop	IP address of the next hop to the destination.	detail, extensive

Table 111: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Next hop type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (locl)—Local address on an interface. • routed multicast (mcrt)—Regular multicast next hop • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp) —Multicast group member. • receive (recv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail, extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail, extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	none detail, extensive
Next-hop interface (Netif)	Interface used to reach the next hop.	none detail, extensive
Alternate forward nh index	Index number of the alternate next hop interface. Seen with multicast option only.	extensive
Next-hop L3 Interface	The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the multicast option.	extensive

Table 111: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Next-hop L2 Interfaces	The next hop layer 2 interfaces. Seen with multicast option only.	extensive

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 extensive
```

```
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                    dscd   50    1
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 route table

List of Syntax

[Syntax on page 3273](#)

[Syntax \(EX Series Switches, QFX Series Switches\) on page 3273](#)

Syntax

```
show route table routing-table-name
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches, QFX Series Switches)

```
show route table routing-table-name
<brief | detail | extensive | terse>
```

Release Information

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.

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display route entries for all routing tables whose names begin with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route table inet** command).

Required Privilege Level

view

RELATED DOCUMENTATION

| [show route summary](#)

List of Sample Output

[show route table bgp.l2vpn.0 on page 3287](#)
[show route table bgp.l3vpn.0 on page 3287](#)
[show route table bgp.l3vpn.0 detail on page 3288](#)
[show route table bgp.rtarget.0 \(When Proxy BGP Route Target Filtering Is Configured\) on page 3290](#)
[show route table bgp.evpn.0 on page 3290](#)
[show route table evpna.evpn.0 on page 3291](#)
[show route table inet.0 on page 3291](#)
[show route table inet.3 on page 3292](#)
[show route table inet.3 protocol ospf on page 3292](#)
[show route table inet6.0 on page 3293](#)
[show route table inet6.3 on page 3293](#)
[show route table inetflow detail on page 3293](#)
[show route table inetflow.0 extensive \(BGP Flowspec Redirect to IP\) on page 3294](#)
[show route table lsdist.0 extensive on page 3296](#)
[show route table l2circuit.0 on page 3298](#)
[show route table lsdist.0 on page 3298](#)
[show route table mpls on page 3299](#)
[show route table mpls extensive on page 3299](#)
[show route table mpls.0 on page 3300](#)
[show route table mpls.0 detail \(PTX Series\) on page 3301](#)
[show route table mpls.0 ccc ge-0/0/1.1004 detail on page 3302](#)
[show route table mpls.0 protocol evpn on page 3303](#)
[show route table mpls.0 protocol ospf on page 3312](#)
[show route table mpls.0 extensive \(PTX Series\) on page 3312](#)
[show route table mpls.0 \(RSVP Route—Transit LSP\) on page 3313](#)
[show route table vpls_1 detail on page 3314](#)
[show route table vpn-a on page 3314](#)
[show route table vpn-a.mdt.0 on page 3315](#)
[show route table VPN-A detail on page 3315](#)
[show route table VPN-AB.inet.0 on page 3316](#)
[show route table VPN_blue.mvpn-inet6.0 on page 3316](#)
[show route table vrf1.mvpn.0 extensive on page 3317](#)
[show route table inetflow detail on page 3318](#)
[show route table bgp.evpn.0 extensive | no-more \(EVPN\) on page 3322](#)
[show route table default-switch.evpn.0 extensive on page 3326](#)
[show route table evpn1.evpn-mcsn on page 3327](#)
[show route table evpn1 \(Multihomed Proxy MAC and IP Address\) on page 3327](#)

Output Fields

Table 112 on page 3275 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

Table 112: 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"> • Pending:protocol-name—List of protocols that have not yet completed graceful restart for this routing table. • Complete—All protocols have restarted for this routing table. <p>For example, if the output shows-</p> <ul style="list-style-type: none"> • LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden) <p>Restart Pending: OSPF LDP VPN</p> <p>This indicates that OSPF, LDP, and VPN protocols did not restart for the LDP.inet.0 routing table.</p> <ul style="list-style-type: none"> • vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden) <p>Restart Complete</p> <p>This indicates that all protocols have restarted for the vpls_1.l2vpn.0 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"> • active (routes that are active) • holddown (routes that are in the pending state before being declared inactive) • hidden (routes that are not used because of a routing policy)

Table 112: 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 entry value is the number of routes for this destination, and the announced 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"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—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. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote. • inclusive multicast Ethernet tag route—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"> • route distinguisher—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI. • Ethernet tag ID—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value. • IP address length—(1 octet) Length of IP address in bits. • originating router's IP address—(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.
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"> • S=0 route 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). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Table 112: show route table Output Fields (*continued*)

Field Name	Field Description
[<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"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
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 Table 113 on page 3282 .
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.

Table 112: show route table Output Fields (*continued*)

Field Name	Field Description
via	<p>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 Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • 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.
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 pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (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 Table 114 on page 3283 .
Local AS	AS number of the local routing devices.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metricn	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.

Table 112: show route table Output Fields (*continued*)

Field Name	Field Description
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 kernel routing table (KRT) for installing the route into the Packet Forwarding Engine, to a resolve tree, a Layer 2 VC, or even a VPN. For example, n-Resolve inet 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"> • n—An index used by Juniper Networks customer support only.
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"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—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. • []—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. • { }—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. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: 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 112: 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"> • Invalid—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. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—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. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Indicates 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, indicates 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, indicates 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.

Table 112: show route table Output Fields (*continued*)

Field Name	Field Description
Communities	Community path attribute for the route. See Table 115 on page 3286 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down .
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).
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.

Table 112: show route table Output Fields (*continued*)

Field Name	Field Description
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 113 on page 3282 describes all possible values for the Next-hop Types output field.

Table 113: 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.
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.

Table 113: Next-hop Types Output Field Values (*continued*)

Next-Hop Type	Description
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 (mcrt)	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 a next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
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 114 on page 3283](#) 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 114: 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.

Table 114: State Output Field Values (*continued*)

Value	Description
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.
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.

Table 114: State Output Field Values (*continued*)

Value	Description
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"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 115 on page 3286 describes the possible values for the Communities output field.

Table 115: 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.
bandwidth: local AS number:link-bandwidth-number	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.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	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.
origin	(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.
route-type-vendor	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 area-number:ospf-route-type:options .
rte-type	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 area-number:ospf-route-type:options .
target	Defines which VPN the route participates in; target has the format 32-bit IP address:16-bit number . For example, 10.19.0.0:100.
unknown IANA	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.

Table 115: Communities Output Field Values (*continued*)

Value	Description
unknown OSPF vendor community	Incoming IANA codes with a value above 0x8000 . This code of the BGP extended community attribute is accepted, but it is not recognized.
evpn-mcast-flags	Identifies the value in the multicast flags extended community and whether snooping is enabled. A value of 0x1 indicates that the route supports IGMP proxy.
evpn-l2-info	Identifies whether Multihomed Proxy MAC and IP Address Route Advertisement is enabled. A value of 0x20 indicates that the proxy bit is set. . Use the show bridge mac-ip-table extensive statement to determine whether the MAC and IP address route was learned locally or from a PE device.

Sample Output

show route table bgp.l2vpn.0

user@host> **show route table bgp.l2vpn.0**

```

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)

10.255.245.12:1:172.16.4.0/8 (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 100.64.12.2 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 100.64.12.2 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 100.64.13.3 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 100.64.13.3 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 100.64.12.2 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 100.64.12.2 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 172.16.5.254 via fxp0.0
10.0.0.1/32        *[Direct/0] 00:51:58
                   > via at-5/3/0.0
10.0.0.2/32        *[Local/0] 00:51:58
                   Local
10.12.12.21/32     *[Local/0] 00:51:57
                   Reject
10.13.13.13/32     *[Direct/0] 00:51:58
                   > via t3-5/2/1.0
10.13.13.14/32     *[Local/0] 00:51:58
                   Local
10.13.13.21/32     *[Local/0] 00:51:58

```



```

Local
10.13.13.22/32    *[Direct/0] 00:33:59
                  > via t3-5/2/0.0
127.0.0.1/32     [Direct/0] 00:51:58
                  > via lo0.0
10.222.5.0/24    *[Direct/0] 00:51:58
                  > via fxp0.0
10.222.5.81/32   *[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

10.0.0.5/32      *[LDP/9] 00:25:43, metric 10, tag 200
                  to 10.2.94.2 via lt-1/2/0.49
                  > to 10.2.3.2 via lt-1/2/0.23

```

show route table inet.3 protocol ospf

user@host> show route table inet.3 protocol ospf

```

inet.3: 9 destinations, 18 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.20/32      [L-OSPF/10] 1d 00:00:56, metric 2
                  > to 10.0.10.70 via lt-1/2/0.14, Push 800020
                  to 10.0.6.60 via lt-1/2/0.12, Push 800020, Push 800030(top)
1.1.1.30/32      [L-OSPF/10] 1d 00:01:01, metric 3
                  > to 10.0.10.70 via lt-1/2/0.14, Push 800030
                  to 10.0.6.60 via lt-1/2/0.12, Push 800030
1.1.1.40/32      [L-OSPF/10] 1d 00:01:01, metric 4
                  > to 10.0.10.70 via lt-1/2/0.14, Push 800040
                  to 10.0.6.60 via lt-1/2/0.12, Push 800040
1.1.1.50/32      [L-OSPF/10] 1d 00:01:01, metric 5
                  > to 10.0.10.70 via lt-1/2/0.14, Push 800050
                  to 10.0.6.60 via lt-1/2/0.12, Push 800050
1.1.1.60/32      [L-OSPF/10] 1d 00:01:01, metric 6
                  > to 10.0.10.70 via lt-1/2/0.14, Push 800060
                  to 10.0.6.60 via lt-1/2/0.12, Pop

```


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: 64502 Peer AS: 64500
              Age: 4
              Task: BGP_64500.10.12.99.5+3792
              Announcement bits (1): 0-Flow
              AS path: 64500 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: 64502
    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 inetflow.0 extensive (BGP Flowspec Redirect to IP)

user@host> show route table inetflow.0 extensive

```

inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
2.2.2.2,*/term:1 (1 entry, 1 announced)
TSI:
KRT in dfwd;
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xb209500 (adv_entry)
Advertised metrics:
Nexthop: 21.1.4.5
Localpref: 100
AS path: [100] I
Communities: redirect-to-ip:21.1.4.5:0
Action(s): accept,count
*Flow Preference: 5
Next hop type: Indirect, Next hop index: 0
Address: 0xa2b931c
Next-hop reference count: 1Next hop:
State: <Active> L
ocal AS: 69
Age: 2
Validation State: unverified
Task: RT Flow
Announcement bits (1): 0-Flow
AS path: I
Communities: redirect-to-ip:21.1.4.5:0

```


user@host> **show route table inetflow.0 extensive**

```
inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
2.2.2.2,*/term:1 (1 entry, 1 announced)
TSI:
KRT in dfwd;
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xb209500 (adv_entry)
Advertised metrics:
Nexthop: 21.1.4.5
Localpref: 100
AS path: [100] I
Communities: redirect-to-nexthop
Action(s): accept,count
*Flow Preference: 5
Next hop type: Indirect, Next hop index: 0
Address: 0xa2b931c
Next-hop reference count: 1
Next hop:
State: <Active>
Local AS: 69
Age: 2
Validation State: unverified
Task: RT Flow
Announcement bits (1): 0-Flow
AS path: I
Communities: redirect-to-nexthop
regress@10.102.178.210> show route table inetflow.0 extensive
inetflow.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
4.4.4.4,*/term:1 (1 entry, 1 announced)
TSI:
KRT in dfwd;
Action(s): accept,count
*BGP Preference: 170/-101
Next hop type: Fictitious, Next hop index: 0
Address: 0xc5e3c30
Next-hop reference count: 3
Next hop: 21.1.4.5
State: <Active Int Ext>
Local AS: 100 Peer AS: 100
Age: 10
Validation State: unverified
Task: BGP_100.1.1.1.1+179
Announcement bits (1): 0-Flow
AS path: I
Communities: redirect-to-nexthop
```



```
Accepted
Localpref: 100
Router ID: 1.1.1.1
```

show route table lsdist.0 extensive

```
user@host> show route table lsdist.0 extensive
```

```
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }/1152
(1 entry, 1 announced)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f378 (adv_entry)
  Advertised metrics:
    Nexthop: Self
    Localpref: 100
    AS path: [4170512532] I
    Communities:
Path NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }
Vector len 4. Val: 0
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Fictitious, Next hop index: 0
    Address: 0x95dfc64
    Next-hop reference count: 9
    State: <Active NotInstall>
    Local AS: 4170512532
    Age: 6:05
    Validation State: unverified
    Task: IS-IS
    Announcement bits (1): 0-BGP_RT_Background
    AS path: I
    IPv4 Router-ids:
      128.220.11.197
    Area membership:
      47 00 05 80 ff f8 00 00 00 01 08 00 01
    SPRING-Capabilities: - SRGB block [Start: 800000,
Range: 256, Flags: 0xc0]
    SPRING-Algorithms:
      - Algo: 0
  LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.3300.5067)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f3cc (adv_entry)
```



```

Advertised metrics:
  Nexthop: Self
  Localpref: 100
  AS path: [4170512532] I
  Communities:
Path LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.33000
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Fictitious, Next hop index: 0
    Address: 0x95dfc64
    Next-hop reference count: 9
    State: <Active NotInstall>
    Local AS: 4170512532
    Age: 6:05
    Validation State: unverified
    Task: IS-IS
    Announcement bits (1): 0-BGP_RT_Background
    AS path: I
    Color: 32768
    Maximum bandwidth: 1000Mbps
    Reservable bandwidth: 1000Mbps
    Unreserved bandwidth by priority:
      0 1000Mbps
      1 1000Mbps
      2 1000Mbps
      3 1000Mbps
      4 1000Mbps
      5 1000Mbps
      6 1000Mbps
      7 1000Mbps
    Metric: 10
    TE Metric: 10
    LAN IPV4 Adj-SID - Label: 299776, Flags: 0x30,
Weight: 0, Nbr: 10.220.1.83

PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 } {
IPv4:128.220.11.197/32 } ISIS-L1:0 }/1152 (1 entry, 1 announced) TSI: Page 0 idx
0, (group ibgp type Internal) Type 1 val 0xa62f43c (adv_entry)
  Advertised metrics:
    Nexthop: Self
    Localpref: 100
    AS path: [4170512532] I
    Communities:

```



```

Path PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 } {
  IPv4:128.220.11.197/32 } ISIS-L1:0 } Vector len 4. Val: 0
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Fictitious, Next hop index: 0
    Address: 0x95dfc64
    Next-hop reference count: 9
    State:<Active NotInstall>
    Local AS: 4170512532
    Age: 6:05
    Validation State: unverified
    Task: IS-IS
    Announcement bits (1): 0-BGP_RT_Background
    AS path: I
      Prefix SID: 67, Flags: 0x40, Algo: 0

```

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 lsdist.0

user@host> show route table lsdist.0


```

lsdist.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

LINK { Local { AS:4 BGP-LS ID:100 IPv4:4.4.4.4 }.{ IPv4:4.4.4.4 } Remote { AS:4
BGP-LS ID:100 IPv4:7.7.7.7 }.{ IPv4:7.7.7.7 } Undefined:0 }/1152
      *[BGP-LS-EPE/170] 00:20:56
      Fictitious
LINK { Local { AS:4 BGP-LS ID:100 IPv4:4.4.4.4 }.{ IPv4:4.4.4.4 IfIndex:339 }
Remote { AS:4 BGP-LS ID:100 IPv4:7.7.7.7 }.{ IPv4:7.7.7.7 } Undefined:0 }/1152
      *[BGP-LS-EPE/170] 00:20:56
      Fictitious
LINK { Local { AS:4 BGP-LS ID:100 IPv4:4.4.4.4 }.{ IPv4:50.1.1.1 } Remote { AS:4
BGP-LS ID:100 IPv4:5.5.5.5 }.{ IPv4:50.1.1.2 } Undefined:0 }/1152
      *[BGP-LS-EPE/170] 00:20:56
      Fictitious

```

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: 18 destinations, 19 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 11:39:56, metric 1
           to table inet.0
0(S=0)     *[MPLS/0] 11:39:56, metric 1
           to table mpls.0
1          *[MPLS/0] 11:39:56, metric 1
           Receive
2          *[MPLS/0] 11:39:56, metric 1
           to table inet6.0
2(S=0)     *[MPLS/0] 11:39:56, metric 1
           to table mpls.0
13         *[MPLS/0] 11:39:56, metric 1
           Receive
303168     *[EVPN/7] 11:00:49, routing-instance pbbn10, route-type
Ingress-MAC, ISID 0
           to table pbbn10.evpn-mac.0
303184     *[EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 1000
           to table pbbn10.evpn-mac.0
           [EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 2000
           to table pbbn10.evpn-mac.0
303264     *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-IM, ISID 1000
           > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303280     *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-IM, ISID 2000
           > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303328     *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
           > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1

```



```

303344          *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
          > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303360          *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:26:88:5f:67:b0
          > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303376          *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:51:51:51:51:51
          > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303392          *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
          > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303408          *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
pbbn10, route-type Egress-MAC, ISID 0
          > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303424          *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC a8:d0:e5:5b:01:c8
          > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303440          *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:52:52:52:52:52
          > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2

```

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: 10.0.0.1 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 ccc ge-0/0/1.1004 detail

user@host>show route table mpls.0 ccc ge-0/0/1.1004 detail

```

mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
ge-0/0/1.1004 (1 entry, 1 announced)
    *EVPN    Preference: 7
            Next hop type: List, Next hop index: 1048577
            Address: 0xdc14770
            Next-hop reference count: 3
            Next hop: ELNH Address 0xd011e30
                Next hop type: Indirect, Next hop index: 0
                Address: 0xd011e30
                Next-hop reference count: 3
                Protocol next hop: 100.100.100.1
                Label operation: Push 301952
                Composite next hop: 0xd011dc0 754 INH Session ID: 0x146
                Indirect next hop: 0xb69a890 1048615 INH Session ID: 0x146
                    Next hop type: Router, Next hop index: 735
                    Address: 0xd00e530
                    Next-hop reference count: 23
                    Next hop: 100.46.1.2 via ge-0/0/5.0
                    Label-switched-path pe4_to_pe1
                    Label operation: Push 300320
                    Label TTL action: prop-ttl
                    Load balance label: Label 300320: None;
                    Label element ptr: 0xd00e580
                    Label parent element ptr: 0x0
                    Label element references: 18
                    Label element child references: 16
                    Label element lsp id: 5
            Next hop: ELNH Address 0xd012070
                Next hop type: Indirect, Next hop index: 0
                Address: 0xd012070
                Next-hop reference count: 3
                Protocol next hop: 100.100.100.2
                Label operation: Push 301888
                Composite next hop: 0xd012000 755 INH Session ID: 0x143

```



```

Indirect next hop: 0xb69a9a0 1048641 INH Session ID: 0x143
  Next hop type: Router, Next hop index: 716
  Address: 0xd00e710
  Next-hop reference count: 23
  Next hop: 100.46.1.2 via ge-0/0/5.0
  Label-switched-path pe4_to_pe2
  Label operation: Push 300304
  Label TTL action: prop-ttl
  Load balance label: Label 300304: None;
  Label element ptr: 0xd00e760
  Label parent element ptr: 0x0
  Label element references: 15
  Label element child references: 13
  Label element lsp id: 6
Next hop: ELNH Address 0xd0121f0, selected
  Next hop type: Indirect, Next hop index: 0
  Address: 0xd0121f0
  Next-hop reference count: 3
  Protocol next hop: 100.100.100.3
  Label operation: Push 301984
  Composite next hop: 0xd012180 756 INH Session ID: 0x145
  Indirect next hop: 0xb69aab0 1048642 INH Session ID: 0x145
    Next hop type: Router, Next hop index: 801
    Address: 0xd010ed0
    Next-hop reference count: 32
    Next hop: 100.46.1.2 via ge-0/0/5.0
    Label-switched-path pe4_to_pe3
    Label operation: Push 300336
    Label TTL action: prop-ttl
    Load balance label: Label 300336: None;
    Label element ptr: 0xd0108c0
    Label parent element ptr: 0x0
    Label element references: 22
    Label element child references: 20
    Label element lsp id: 7
State: < Active Int >
Age: 2:06:50
Validation State: unverified
Task: evpn global task
Announcement bits (1): 1-KRT
AS path: I

```

show route table mpls.0 protocol evpn

user@host>**show route table mpls.0 protocol evpn**


```

mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299872          *[EVPN/7] 02:30:58, routing-instance mhevpn, route-type
Ingress-IM, vlan-id 10
                to table mhevpn.evpn-mac.0
300016          *[EVPN/7] 02:30:38, routing-instance VS-1, route-type Ingress-IM,
vlan-id 110
                to table VS-1.evpn-mac.0
300032          *[EVPN/7] 02:30:38, routing-instance VS-1, route-type Ingress-IM,
vlan-id 120
                to table VS-1.evpn-mac.0
300048          *[EVPN/7] 02:30:38, routing-instance VS-1, route-type Ingress-IM,
vlan-id 130
                to table VS-1.evpn-mac.0
300064          *[EVPN/7] 02:30:38, routing-instance VS-2, route-type Ingress-IM,
vlan-id 210
                to table VS-2.evpn-mac.0
300080          *[EVPN/7] 02:30:38, routing-instance VS-2, route-type Ingress-IM,
vlan-id 220
                to table VS-2.evpn-mac.0
300096          *[EVPN/7] 02:30:38, routing-instance VS-2, route-type Ingress-IM,
vlan-id 230
                to table VS-2.evpn-mac.0
300112          *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300128          *[EVPN/7] 02:29:22, routing-instance mhevpn, route-type
Ingress-Aliasing
                to table mhevpn.evpn-mac.0
300144          *[EVPN/7] 02:27:06, routing-instance VS-1, route-type Egress-MAC,
ESI 00:44:44:44:44:44:44:44:44:44
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300160          *[EVPN/7] 02:29:22, routing-instance VS-1, route-type
Ingress-Aliasing
                to table VS-1.evpn-mac.0
300176          *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
ESI 00:44:44:44:44:44:44:44:44:44
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300192          *[EVPN/7] 02:29:22, routing-instance VS-2, route-type
Ingress-Aliasing
                to table VS-2.evpn-mac.0
300208          *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-IM, vlan-id 120

```



```

> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300224      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
mhevpn, route-type Egress-IM, vlan-id 10
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300240      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-IM, vlan-id 110
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300256      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-IM, vlan-id 130
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300272      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-2, route-type Egress-IM, vlan-id 210
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300288      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-2, route-type Egress-IM, vlan-id 220
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300304      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-2, route-type Egress-IM, vlan-id 230
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300320      *[EVPN/7] 02:27:06, routing-instance VS-1, route-type Egress-MAC,
ESI 00:11:11:11:11:11:11:11:11:11
to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300336      *[EVPN/7] 02:27:06, routing-instance VS-1, route-type Egress-MAC,
ESI 00:33:33:33:33:33:33:33:33:33
to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300368      *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
ESI 00:33:33:33:33:33:33:33:33:33
to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300384      *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
ESI 00:11:11:11:11:11:11:11:11:11
to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300416      *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type

```



```

Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300432    *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300480    *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300496    *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-2, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300560    *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-1, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300592    *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
VS-2, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300608    *[EVPN/7] 02:29:23
    > via ge-0/0/1.1001, Pop
300624    *[EVPN/7] 02:29:23
    > via ge-0/0/1.2001, Pop
301232    *[EVPN/7] 02:29:17
    > via ge-0/0/1.1002, Pop
301296    *[EVPN/7] 02:29:10
    > via ge-0/0/1.1003, Pop
301312    *[EVPN/7] 02:27:06
    > via ae10.2003, Pop
    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301360    *[EVPN/7] 02:29:01
    > via ge-0/0/1.1004, Pop
301408    *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
vpws1004, route-type Egress, vlan-id 2004
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301456    *[EVPN/7] 02:27:06
    > via ae10.1010, Pop
    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301552    *[EVPN/7] 02:27:07, routing-instance VS-1, route-type Egress-MAC,
ESI 00:22:22:22:22:22:22:22:22:22

```



```

> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301568      *[EVPN/7] 02:27:07, routing-instance VS-2, route-type Egress-MAC,
ESI 00:22:22:22:22:22:22:22:22
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301648      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
vpws1010, route-type Egress, vlan-id 2010
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301664      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
mhevpn, route-type Egress-MAC
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301680      *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
mhevpn, route-type Egress-MAC
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301696      *[EVPN/7] 02:27:07, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301712      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-MAC
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301728      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-MAC
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301744      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-IM, vlan-id 230
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301760      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
vpws1010, route-type Egress, vlan-id 2010
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301776      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
mhevpn, route-type Egress-MAC
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301792      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-IM, vlan-id 130
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301808      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
vpws1004, route-type Egress, vlan-id 2004
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301824      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
mhevpn, route-type Egress-IM, vlan-id 10
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301840      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
vpws1002, route-type Egress, vlan-id 2002
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301856      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance

```



```

vpws1003, route-type Egress, vlan-id 2003
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301872      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
vpws1003, route-type Egress Protection, vlan-id 2003
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301888      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
vpws1010, route-type Egress Protection, vlan-id 1010
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301904      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-IM, vlan-id 220
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301920      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-IM, vlan-id 210
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301936      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-IM, vlan-id 230
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301952      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-SH, vlan-id 230
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301968      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-IM, vlan-id 220
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301984      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-SH, vlan-id 220
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302000      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-IM, vlan-id 210
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302016      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-SH, vlan-id 210
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302032      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302048      *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-2, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302064      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302080      *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-2, route-type Egress-MAC
    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

```



```

302096          *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-MAC
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302112          *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-MAC
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302128          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-MAC
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302144          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-MAC
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302160          *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-IM, vlan-id 120
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302176          *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
VS-1, route-type Egress-IM, vlan-id 110
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302192          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-IM, vlan-id 130
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302208          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-SH, vlan-id 130
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302224          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-IM, vlan-id 120
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302240          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-SH, vlan-id 120
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302256          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-IM, vlan-id 110
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302272          *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
VS-1, route-type Egress-SH, vlan-id 110
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302288          *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
mhevnpn, route-type Egress-MAC
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302304          *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
mhevnpn, route-type Egress-MAC
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302320          *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
mhevnpn, route-type Egress-MAC

```



```

> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302336      *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
mhevpn, route-type Egress-MAC
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302352      *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
vpws1004, route-type Egress, vlan-id 2004
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302368      *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
mhevpn, route-type Egress-IM, vlan-id 10
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302384      *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
mhevpn, route-type Egress-SH, vlan-id 10
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302400      *[EVPN/7] 02:26:21
> via ge-0/0/1.3001, Pop
302432      *[EVPN/7] 02:26:21, remote-pe 100.100.100.3, routing-instance
vpws3001, route-type Egress, vlan-id 40000
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302448      *[EVPN/7] 02:26:21, remote-pe 100.100.100.1, routing-instance
vpws3001, route-type Egress, vlan-id 40000
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302464      *[EVPN/7] 02:26:20, remote-pe 100.100.100.2, routing-instance
vpws3001, route-type Egress, vlan-id 40000
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302480      *[EVPN/7] 02:26:14
> via ge-0/0/1.3016, Pop
302512      *[EVPN/7] 02:26:14, remote-pe 100.100.100.1, routing-instance
vpws3016, route-type Egress, vlan-id 40016
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302528      *[EVPN/7] 02:26:14, remote-pe 100.100.100.2, routing-instance
vpws3016, route-type Egress, vlan-id 40016
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302560      *[EVPN/7] 02:26:06
> via ae10.3011, Pop
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302592      *[EVPN/7] 02:26:07, remote-pe 100.100.100.1, routing-instance
vpws3011, route-type Egress, vlan-id 401100
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302608      *[EVPN/7] 02:26:07, remote-pe 100.100.100.2, routing-instance
vpws3011, route-type Egress, vlan-id 401100
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302624      *[EVPN/7] 02:26:07, remote-pe 100.100.100.3, routing-instance
vpws3011, route-type Egress Protection, vlan-id 301100
> to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

```



```

302656          *[EVPN/7] 02:25:59
                > via ae10.3006, Pop
                to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302688          *[EVPN/7] 02:26:00, remote-pe 100.100.100.2, routing-instance
vpws3006, route-type Egress, vlan-id 400600
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302704          *[EVPN/7] 02:26:00, remote-pe 100.100.100.1, routing-instance
vpws3006, route-type Egress, vlan-id 400600
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302720          *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
vpws3006, route-type Egress, vlan-id 400600
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302736          *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
vpws3006, route-type Egress Protection, vlan-id 300600
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.1001    *[EVPN/7] 02:29:23
                > via ge-0/0/1.2001
ge-0/0/1.2001    *[EVPN/7] 02:29:23
                > via ge-0/0/1.1001
ge-0/0/1.1002    *[EVPN/7] 02:27:06
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.2003        *[EVPN/7] 02:29:10
                > via ge-0/0/1.1003
ge-0/0/1.1003    *[EVPN/7] 02:27:06
                to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

                > via ae10.2003
ge-0/0/1.1004    *[EVPN/7] 02:27:06
                to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.1010        *[EVPN/7] 02:27:06
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ge-0/0/1.3001    *[EVPN/7] 02:26:20
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.3016    *[EVPN/7] 02:26:13
                > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3011        *[EVPN/7] 02:26:06

```



```

ae10.3006          > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
                   *[EVPN/7] 02:25:59
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                   to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                   to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

```

show route table mpls.0 protocol ospf

user@host> show route table mpls.0 protocol ospf

```

mpls.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299952          *[L-OSPF/10] 23:59:42, metric 0
                > to 10.0.10.70 via lt-1/2/0.14, Pop
                to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299952(S=0)     *[L-OSPF/10] 23:59:42, metric 0
                > to 10.0.10.70 via lt-1/2/0.14, Pop
                to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299968          *[L-OSPF/10] 23:59:48, metric 0
                > to 10.0.6.60 via lt-1/2/0.12, Pop

```

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: 10.0.0.1 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: 10.0.0.1 via ge-0/0/1.0
        Session Id: 0x1
    10.255.255.1/32 Originating RIB: inet.3
        Metric: 1                      Node path count: 1
        Forwarding nexthops: 1
            Nexthop: 10.0.0.1 via ge-0/0/1.0

```

show route table mpls.0 (RSVP Route—Transit 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 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p

```



```

300352(S=0)      *[RSVP/7/1] 00:08:00, metric 1
                  > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384           *[RSVP/7/2] 00:05:20, metric 1
                  > to 10.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)      *[RSVP/7/2] 00:05:20, metric 1
                  > to 10.64.1.106 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

172.16.1.11:1000:1:1/96 (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.l2vpn.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.l3vpn.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

1:10.255.2.202:65536:10.255.2.202/432
    *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
        AS path: I
        > via so-0/1/3.0
1:10.255.2.203:65536:10.255.2.203/432
    *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
        AS path: I
        > via so-0/1/0.0
1:10.255.2.204:65536:10.255.2.204/432
    *[MVPN/70] 00:57:23, metric2 1
        Indirect
5:10.255.2.202:65536:128:::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
6:10.255.2.203:65536:64500:128:::10.12.53.12:128:ffff::1/432
    *[PIM/105] 00:02:37
        Multicast (IPv6)
7:10.255.2.202:65536:64500:128:::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 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: 64502 Peer AS: 64500
              Age: 4
              Task: BGP_64500.10.12.99.5+3792
              Announcement bits (1): 0-Flow
              AS path: 64500 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: 64502
              Age: 6:30
              Task: RT Flow
              Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
              AS path: I
              Communities: 1:1
```

```
user@host> 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

10.1.1.2:100:10.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
10.1.1.4:100:10.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
10.1.1.2:100:1:0/96 MH
```



```

          *[VPLS/170] 1d 03:11:03, metric2 1
          Indirect
10.1.1.4:100:1:0/96 MH
          *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
          AS path: I, validation-state: unverified
          > via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.2:10.1.1.4/176
          *[VPLS/7] 1d 03:11:02, metric2 1
          > via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.4:10.1.1.2/176
          *[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)
10.0.0.0/32 (3 entries, 1 announced)
    State: <OnList CalcForwarding>
TSI:
KRT in-kernel 10.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 10.0.0.0 from 10.3.0.0 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: 10.2.0.0
            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_172.16.2.0.0+34549
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 10.2.0.0
Primary Routing Table bgp.l3vpn.0
Composite next hops: 1
    Protocol next hop: 10.2.0.0 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: 10.3.0.0
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: 10.3.0.0
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_172.16.3.0.0+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: 10.3.0.0
Primary Routing Table bgp.l3vpn.0
Composite next hops: 1
    Protocol next hop: 10.3.0.0 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
    10.3.0.0/32 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: 10.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

    Protocol next hop: 10.3.0.0
    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)

user@host> **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: 10.2.3.4
            Protocol next hop: 10.2.3.4
            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: 10.2.3.4
            Secondary Tables: default-switch.evpn.0
            Indirect next hops: 1
              Protocol next hop: 10.2.3.4 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: 10.2.3.4
        Protocol next hop: 10.2.3.4
        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: 00:00:00:00:00:00:00:00:00:00:00
    Localpref: 100
    Router ID: 10.2.3.4
    Secondary Tables: default-switch.evpn.0
    Indirect next hops: 1
      Protocol next hop: 10.2.3.4 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
      10.2.3.4/32 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: 10.2.3.4

```



```

        Protocol next hop: 10.2.3.4
        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: 00:00:00:00:00:00:00:00:00:00
        Localpref: 100
        Router ID: 10.2.3.4
        Secondary Tables: default-switch.evpn.0
        Indirect next hops: 1
            Protocol next hop: 10.2.3.4 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
            10.2.3.4/32 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: 10.2.3.4
        Protocol next hop: 10.2.3.4
        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: 10.2.3.4
        Secondary Tables: default-switch.evpn.0
        Indirect next hops: 1
            Protocol next hop: 10.2.3.4 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
        10.2.3.4/32 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: 10.2.3.4
              Protocol next hop: 10.2.3.4
              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: 10.2.3.4
    Secondary Tables: default-switch.evpn.0
    Indirect next hops: 1
        Protocol next hop: 10.2.3.4 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
        10.2.3.4/32 Originating RIB: inet.0
            Metric: 1                      Node path count: 1
            Forwarding nexthops: 2
                Nexthop: 10.92.78.102 via em0.0

3:1000:10::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: 10.2.3.4
        Protocol next hop: 10.2.3.4
        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: 10.2.3.4
        Secondary Tables: default-switch.evpn.0
        Indirect next hops: 1
            Protocol next hop: 10.2.3.4 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
            10.2.3.4/32 Originating RIB: inet.0
                Metric: 1 Node path count: 1
                Forwarding nexthops: 2
                    Nexthop: 10.92.78.102 via em0.0

```

show route table default-switch.evpn.0 extensive

The following shows the partial output listing for the EVPN VNI table.

user@host> **show route table default-switch.evpn.0 extensive**

```

3:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 1 announced)
    *BGP Preference: 170/-101
        Route Distinguisher: 10.255.0.1:00
        PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
        Next hop type: Indirect, Next hop index: 0
        Address: 0xcebfad0
        Next-hop reference count: 26
        Source: 10.255.0.1
        Protocol next hop: 10.255.0.1
        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
        State: <Secondary Active Int Ext>
    Local AS: 100 Peer AS: 100
        Age: 1:35:30 Metric2: 2
        Validation State: unverified

```



```

Task: BGP_100.10.255.0.1
Announcement bits (1): 0-default-switch-evpn
AS path: I
Communities: target:100:100 encapsulation:vxlan (0x8)
evpn-mcast-flags:0x1:snooping-enabled
. . .

```

show route table evpn1.evpn-mcsn

The following shows the output listing for the multicast information used by the rpd and mcsnoopd.

user@host> **show route table default-switch.evpn-mcsn.1**

```

default-switch.evpn-mcsn.1: 9 destinations, 9 routes (9 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

0.14,0.0,0.0/48      *[Multicast/180] 00:01:02
                    to 1.1.1.1 via vtep.32770
                    to 1.2.2.2 via vtep.32771
                    to 1.6.6.6 via vtep.32769
                    to 1.3.3.3 via vtep.32772
0.14,0.0,0.0,224.0.0.0/52*[Multicast/180] 00:01:02
                    to 1.1.1.1 via vtep.32770
                    to 1.2.2.2 via vtep.32771
                    to 1.6.6.6 via vtep.32769
0.14,0.0,0.0,225.1.1.1/80*[Multicast/180] 00:00:06
                    to 1.1.1.1 via vtep.32770
                    to 1.2.2.2 via vtep.32771
                    to 1.6.6.6 via vtep.32769
                    to 1.3.3.3 via vtep.32772

```

show route table evpn1 (Multihomed Proxy MAC and IP Address)

The following shows a partial output listing for an EVPN instance. This indicates when Multihomed Proxy MAC and IP Address Route Advertisement is enabled.

user@host> **show route table evpn-1**

```

2:666:11010003::1002::00:00:00:00:00:02::102.1.1.2/304 MAC/IP (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group vteps type Internal) Type 1 val 0xb20eb10 (adv_entry)
  Advertised metrics:

```



```
Nexthop: 103.1.1.1
Localpref: 100
AS path: [666] I
Communities: target:666:1002 evpn-l2-info:0x20:proxy (mtu 0)
Path 2:666:11010003::1002::00:00:00:00:00:02::102.1.1.2 Vector len 4. Val: 0
  *EVPN   Preference: 170
          Next hop type: Indirect, Next hop index: 0
          Address: 0xc3a9cf0
          Next-hop reference count: 36
          Protocol next hop: 103.1.1.1
          Indirect next hop: 0x0 - INH Session ID: 0x0
          State: <Active Int Ext>
```


show ted database

List of Syntax

[Syntax on page 3329](#)

[Syntax \(EX Series Switches\) on page 3329](#)

Syntax

```
show ted database
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
<system-name>
<topology-id topology bgp-ls-epe>
```

Syntax (EX Series Switches)

```
show ted database
<brief | detail | extensive>
<system-name>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

topology-id *topology* option added in Junos OS Release 17.4R1 for MX Series and PTX Series.

Description

Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.

Options

none—Display standard information about all entries in the traffic engineering database.

brief | detail | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, 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.

system-name—(Optional) Display traffic engineering database information for a particular system.

topology-id *topology*— Display the topology information. By default, traffic engineering topology information is displayed.

Required Privilege Level

view

List of Sample Output

[show ted database brief on page 3333](#)

[show ted database detail on page 3334](#)

[show ted database extensive on page 3335](#)

[show ted database topology-id igp on page 3338](#)

[show ted database topology-id bgp-ls-epe extensive on page 3339](#)

Output Fields

[Table 116 on page 3330](#) describes the output fields for the **show ted database** command. Output fields are listed in the approximate order in which they appear.

Table 116: 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
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.	brief
NodeID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	extensive
Type	Type of node. It can be either Rtr (router) or Net (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: <ul style="list-style-type: none"> • IS-IS(1)—IS-IS Level 1. • IS-IS(2)—IS-IS Level 2. • OSPF (area-number)—OSPF from the specified area. 	All levels
To	Address on the far end of a link.	detail extensive

Table 116: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local	Address of the local interface being used to reach the remote node.	detail extensive
Remote	Address of the interface on the remote node.	detail extensive
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail extensive
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail extensive
Metric	Configured traffic engineering metric.	extensive
IGP metric	Configured interior gateway protocol metric.	extensive
Static BW	Total interface bandwidth in bps.	extensive
Reservable bandwidth	Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the subscription statement when configuring RSVP.	extensive
Available BW [priority]	(Must include diffserv-te statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.	extensive
Diffserv-TE BW Model	Bandwidth constraint model used by the LSPs.	extensive
Available BW [TE-class]	(Must include the diffserv-te statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.	extensive
Static BW [CT-class]	Total interface bandwidth used by an MPLS traffic class, in bps.	extensive

Table 116: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Interface Switching Capability Descriptor (n)	<p>Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. <i>n</i> is the index number.</p> <ul style="list-style-type: none"> • Switching type—Type of switching to be performed on a particular link: <ul style="list-style-type: none"> • PSC-1—Packet switch-capable 1 • PSC-2—Packet switch-capable 2 • PSC-3—Packet switch-capable 3 • PSC-4—Packet switch-capable 4 • L2SC—Layer-2-switch-capable • TDM—Time-division-multiplexing-capable • LSC—Lambda switch-capable • FSC—Fiber switch-capable • Encoding type—Encoding of the LSP being requested: <ul style="list-style-type: none"> • Packet • Ethernet • ANSI/ETSI PDH • Reserved • SDH /SONET • Digital Wrapper • Lambda (photonic) • Fiber • FiberSDH/SONET • Maximum LSP BW [priority] bps—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <ul style="list-style-type: none"> • [n]—Priority level. The range is from 0 (high) through 7 (low). • n Mbps—Amount of the maximum bandwidth. • Minimum LSP BW—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. Minimum LSP BW is displayed only when switching type is PSC-1 or TDM. • Interface MTU—Displayed only when switching type is TDM. • Interface supports standard SONET/SDH—Displayed only when switching type is TDM. 	extensive

Sample Output

show ted database brief

user@host> show ted database brief

```
TED database: 12 ISIS nodes 0 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---   3178    2     0
Router-B.00                      ---   3152    2     0
Router-B.02                      Net    802    0     2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---   3126    2     0
Router-C.02                      Net    38     0     2 IS-IS(2)
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---   3144    2     0
Router-D.02                      Net    723    0     2 IS-IS(2)
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net    607    0     2 IS-IS(2)
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---   3178    2     0
Router-E.02                      Net    131    0     2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---   3153    2     0
Router-F.02                      Net    769    0     2 IS-IS(2)
```



```

To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0

```

show ted database detail

```

TED database: 12 ISIS nodes 0 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---   2913     2     0
Router-B.00                      ---   2887     2     0
Router-B.02                      Net    537     0     2 IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---   2861     2     0
Router-C.02                      Net    597     0     2 IS-IS(2)
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---   2879     2     0
Router-D.02                      Net    458     0     2 IS-IS(2)
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net    342     0     2 IS-IS(2)
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---   2913     2     0
Router-E.02                      Net    640     0     2 IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```



```

ID                               Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---   2888    2      0
Router-F.02                      Net    504    0      2 IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0

```

show ted database extensive

user@host> show ted database extensive

```

TED database: 12 ISIS nodes 0 INET nodes
NodeID: Router-A.00
    Type: ---, Age: 3067 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.00
    Type: ---, Age: 3041 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.02
    Type: Net, Age: 691 secs, LinkIn: 0, LinkOut: 2
    Protocol: IS-IS(2)
        To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
            Local interface index: 0, Remote interface index: 0
            Metric: 0
            IGP metric: 10
            Interface Switching Capability Descriptor(1):
                Switching type: Packet
                Encoding type: Packet
                Maximum LSP BW [priority] bps:
                    [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
                    [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
            To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
                Local interface index: 0, Remote interface index: 0
                Metric: 0
                IGP metric: 20
                Interface Switching Capability Descriptor(1):
                    Switching type: Packet
                    Encoding type: Packet
                    Maximum LSP BW [priority] bps:
                        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
                        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-C.00
    Type: ---, Age: 3015 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-C.02
    Type: Net, Age: 751 secs, LinkIn: 0, LinkOut: 2

```



```

Protocol: IS-IS(2)
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    IGP metric: 10
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    IGP metric: 10      Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-D.00
  Type: ---, Age: 3034 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-D.02
  Type: Net, Age: 613 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    IGP metric: 10
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    IGP metric: 10
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps

```



```

                [4] 0bps          [5] 0bps          [6] 0bps          [7] 0bps
NodeID: Router-D.03
Type: Net, Age: 497 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    IGP metric: 10
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps          [1] 0bps          [2] 0bps          [3] 0bps
        [4] 0bps          [5] 0bps          [6] 0bps          [7] 0bps
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    IGP metric: 10
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps          [1] 0bps          [2] 0bps          [3] 0bps
        [4] 0bps          [5] 0bps          [6] 0bps          [7] 0bps
NodeID: Router-E.00
Type: ---, Age: 3068 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-E.02
Type: Net, Age: 21 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    Interface Switching Capability Descriptor(1):
      Switching type: Packet
      Encoding type: Packet
      Maximum LSP BW [priority] bps:
        [0] 0bps          [1] 0bps          [2] 0bps          [3] 0bps
        [4] 0bps          [5] 0bps          [6] 0bps          [7] 0bps
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    Metric: 0
    IGP metric: 10
    Interface Switching Capability Descriptor(1):
      Switching type: Packet

```



```

        Encoding type: Packet
        Maximum LSP BW [priority] bps:
            [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
            [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-F.00
    Type: ---, Age: 3043 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-F.02
    Type: Net, Age: 659 secs, LinkIn: 0, LinkOut: 2
    Protocol: IS-IS(2)
        To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        IGP metric: 10
        Interface Switching Capability Descriptor(1):
            Switching type: Packet
            Encoding type: Packet
            Maximum LSP BW [priority] bps:
                [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
                [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
        To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
        Local interface index: 0, Remote interface index: 0
        Metric: 0
        IGP metric: 10      Interface Switching Capability Descriptor(1):
            Switching type: Packet
            Encoding type: Packet
            Maximum LSP BW [priority] bps:
                [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
                [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps

```

show ted database topology-id igp

user@host> show ted database topology-id igp

```

TED database: 3 ISIS nodes 3 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router A.00(128.220.1.2)         Rtr   193     2     2 IS-IS(2)
    To: Router B.00(128.220.18.198), Local: 2.3.0.2, Remote: 2.3.0.1
        Local interface index: 334, Remote interface index: 336
    To: Router B.00(128.220.18.198), Local: 2.3.1.2, Remote: 2.3.1.1
        Local interface index: 333, Remote interface index: 335
ID                               Type Age(s) LnkIn LnkOut Protocol
Router C.00(128.220.1.52)         Rtr   193     2     2 IS-IS(2)
    To: Router B.00(128.220.18.198), Local: 1.2.0.1, Remote: 1.2.0.2
        Local interface index: 335, Remote interface index: 334

```



```

To: Router B.00(128.220.18.198), Local: 1.2.1.1, Remote: 1.2.1.2
  Local interface index: 334, Remote interface index: 333
ID                                     Type Age(s) LnkIn LnkOut Protocol
Router B.00(128.220.18.198)    Rtr   193     4     4 IS-IS(2)
  To: Router A.00(128.220.1.2), Local: 2.3.0.1, Remote: 2.3.0.2
    Local interface index: 336, Remote interface index: 334
  To: Router A.00(128.220.1.2), Local: 2.3.1.1, Remote: 2.3.1.2
    Local interface index: 335, Remote interface index: 333
  To: Router C.00(128.220.1.52), Local: 1.2.0.2, Remote: 1.2.0.1
    Local interface index: 334, Remote interface index: 335
  To: Router C.00(128.220.1.52), Local: 1.2.1.2, Remote: 1.2.1.1
    Local interface index: 333, Remote interface index: 334

```

show ted database topology-id bgp-ls-epe extensive

user@host> show ted database topology-id bgp-ls-epe extensive

```

TED database: 0 ISIS nodes 3 INET nodes
NodeID: 4.4.4.4                                     << DUT router-id
  Type: Rtr, Age: 270 secs, LinkIn: 0, LinkOut: 3
  Protocol: BGP-LS-EPE(0)                             << Protocol
    To: 5.5.5.5, Local: 50.1.1.1, Remote: 50.1.1.2 << Peer router-id and local and
remote interface used for BGP session.
      Local interface index: 0, Remote interface index: 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
      BGP-Peer-SID:                                     << BGP-Peer-SID Information
        SID: 1000007, Type: Node-SID Flags: 0x30, Weight: 0 << BGP-Node-SID
        SID: 1000002, Type: Set-SID Flags: 0x30, Weight: 0 << BGP-Set-SID
    To: 7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7
      Local interface index: 0, Remote interface index: 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
      BGP-Peer-SID:
        SID: 1000006, Type: Node-SID Flags: 0x30, Weight: 0 << BGP-Node-SID
    To: 7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7

```



```

Local interface index: 339, Remote interface index: 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
  BGP-Peer-SID:
    SID: 1000005, Type: Adj-SID Flags: 0x30, Weight: 0 << BGP-Adj-SID
NodeID: 5.5.5.5
  Type: Rtr, Age: 270 secs, LinkIn: 1, LinkOut: 0
  Protocol: BGP-LS-EPE(0)
NodeID: 7.7.7.7
  Type: Rtr, Age: 270 secs, LinkIn: 2, LinkOut: 0
  Protocol: BGP-LS-EPE(0)

```


show ted link

List of Syntax

[Syntax on page 3341](#)

[Syntax \(EX Series Switches\) on page 3341](#)

Syntax

```
show ted link
<brief | detail>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show ted link
<brief | detail>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.

Options

none—Display standard information about traffic engineering database link information.

brief | detail—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, 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.

Required Privilege Level

view

List of Sample Output

[show ted link brief on page 3343](#)

[show ted link detail on page 3343](#)

[show ted link topology-id bgp-ls-epe detail on page 3344](#)

Output Fields

Table 117 on page 3342 describes the output fields for the **show ted link** command. Output fields are listed in the approximate order in which they appear.

Table 117: show ted link Output Fields

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	brief
-->ID	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	brief
<i>hostname</i>	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	detail
<i>hostname</i>	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	detail
Local Path	Number of paths CSPF on the local routing device has placed on the link.	All levels
Metric	Configured traffic engineering metric.	extensive
IGP metric	Configured interior gateway protocol metric.	detail
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels
Local	Address of the local interface being used to reach the remote node.	detail extensive
Remote	Address of the interface on the remote node.	detail extensive
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail

Sample Output

show ted link brief

```
user@host> show ted link brief
```

ID	->ID	LocalPath	LocalBW
Router-B.02	Router-A.00		0 0bps
Router-B.02	Router-B.00		0 0bps
Router-C.02	Router-B.00		0 0bps
Router-C.02	Router-C.00		0 0bps
Router-D.02	Router-F.00		0 0bps
Router-D.02	Router-D.00		0 0bps
Router-D.03	Router-D.00		0 0bps
Router-D.03	Router-C.00		0 0bps
Router-E.02	Router-A.00		0 0bps
Router-E.02	Router-E.00		0 0bps
Router-F.02	Router-E.00		0 0bps
Router-F.02	Router-F.00		0 0bps

show ted link detail

```
user@host> show ted link detail
```

```
Router-B.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
```



```

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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
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: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps

```

show ted link topology-id bgp-ls-epe detail

```
user@host> show ted link topology-id bgp-ls-epe detail
```



```
4.4.4.4->5.5.5.5, Local: 50.1.1.1, Remote: 50.1.1.2
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
  SID: 1000007 Type: Node-SID Flags: 0x30 Weight: 0
  SID: 1000002 Type: Set-SID Flags: 0x30 Weight: 0
4.4.4.4->7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
  SID: 1000006 Type: Node-SID Flags: 0x30 Weight: 0
4.4.4.4->7.7.7.7, Local: 4.4.4.4, Remote: 7.7.7.7
  Local interface index: 339, Remote interface index: 0
  LocalPath: 0, AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
  SID: 1000005 Type: Adj-SID Flags: 0x30 Weight: 0
```


show ted protocol

List of Syntax

[Syntax on page 3346](#)

[Syntax \(EX Series Switches\) on page 3346](#)

Syntax

```
show ted protocol
<brief | detail>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show ted protocol
<brief | detail>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance *instance-name* option added in Junos OS Release 15.1.

Description

Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.

Options

none—Display standard information about the protocols from which the traffic engineering database learned about its nodes.

brief | detail—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, 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.

Required Privilege Level

view

List of Sample Output

[show ted protocol on page 3347](#)

Output Fields

[Table 118 on page 3347](#) describes the output fields for the **show ted protocol** command. Output fields are listed in the approximate order in which they appear.

Table 118: show ted protocol Output Fields

Field Name	Field Description
Protocol name	Protocol that reported the node information: <ul style="list-style-type: none"> • IS-IS(1)—IS-IS Level 1. • IS-IS(2)—IS-IS Level 2. • OSPF (<i>area-number</i>)—OSPF from the specified area.
Credibility	If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.
Self node	Address the protocol uses as the local address.

Sample Output

show ted protocol

```
user@host> show ted protocol
```

```

Protocol name      Credibility  Self node
IS-IS(2)           2  (highest) corriedale.00(123.456.1.11)
IS-IS(1)           1              corriedale.00(123.456.1.11)
```

```
user@host> show ted protocol topology-id bgp-ls-epe detail
```

```

Protocol name      Credibility  Self node
BGP-LS-EPE(0)     342          200.0.0.4
```


traceroute mpls bgp

Syntax

```
traceroute mpls bgp fec
<destination destination-address>
<detail>
<exp exp>
<fanout fanout-number>
<logical-system logical system-name>
<no-resolve>
<paths paths-number>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source source-address>
<ttl value>
<wait seconds>
```

Release Information

Command introduced in Junos OS Release 14.2.

Description

Trace route to a remote host for an MPLS label-switched path (LSP) signaled by the Border Gateway Protocol (BGP). Use **traceroute mpls bgp** as a debugging tool to locate MPLS BGP forwarding issues in a network. (Currently supported for IPv4 packets only.)

To use the **traceroute mpls bgp** command, make sure you have BGP routes with MPLS labels.

NOTE: The **traceroute mpls bgp fec** command only supports single paths.

Options

fec—Specify the IP address and optional prefix of the forwarding equivalence class (FEC). Suppose you are at PE1, use would want to use the IP address of PE2 to trace the BGP path to that router.

destination *destination-address*—(Optional) Specify the destination address to use when sending probes.

detail—(Optional) Display detailed output.

exp *exp*—(Optional) Specify the class of service to use when sending probes.

Range: 0 through 7

Default: 7

fanout *fanout-number*—(Optional) Specify the maximum number of next hops to search per node.

Range: 1 through 16

Default: 16

logical-system *logical-system-name*—(Optional) Specify the name of the logical system for the traceroute attempt.

no-resolve—(Optional) Specify not to resolve the hostname that corresponds to the IP address.

paths *paths-number*—(Optional) Specify the number of paths to search.

Range: 1 through 255

Default: 16

pipe-mode—(Optional) Specify to trace only the outermost FEC.

retries *retries-number*—(Optional) Specify the number of times to resend probe values.

Range: 1 through 9

Default: 3

routing-instance *routing-instance-name*—(Optional) Specify the name of the routing instance for the trace route attempt.

source *source-address*—(Optional) Specify the source address of the outgoing traceroute packets.

ttl *value*—(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds.

Range: 1 through 125

Default: 64

wait *seconds*—(Optional) Specify the number of seconds to wait before resending a probe.

Range: 5 through 15

Default: 10

Required Privilege Level

network

RELATED DOCUMENTATION

| [ping mpls bgp](#) | [3125](#)

List of Sample Output

[traceroute mpls bgp on page 3351](#)

[traceroute mpls bgp detail on page 3351](#)

Output Fields

Table 119 on page 3350 describes the output fields for the **traceroute mpls bgp fec** command and the **traceroute mpls bgp fec detail** command. Output fields are listed in the approximate order in which they appear.

Table 119: traceroute mpls bgp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls bgp fec command.	All levels
ttl	Time to live value of the labeled packet.	None
Label	Outgoing label used for forwarding the packet along the label-switched paths.	None
Protocol	Signaling protocol used. For this command, it is BGP.	None
Address	Address of the next hop.	None
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null .	None
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	None
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null .	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail
Response time	Time for the echo request to reach the receiver.	detail
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none .	detail
Label Stack	Label stack used to forward the packet.	detail

Sample Output

traceroute mpls bgp

user@host> traceroute mpls bgp fec

```

    Probe options: retries 3, exp 7
  ttl Label  Protocol Address  Previous Hop  Probe Status Fec-Stack-Sent Fec-Change
  1   299824  LDP      81.1.2.2    (null)       Success     LDP, BGP      PUSH-RSVP
  2   299825  RSVP     81.2.3.3    81.1.2.2    Success     RSVP, LDP, BGP (null)
  3   299826  RSVP     81.3.4.4    81.2.3.3    Egress      RSVP, LDP, BGP POP-RSVP
  3   299826  LDP      81.3.4.4    81.2.3.3    Success     LDP, BGP      (null)
  4   299827  LDP      81.4.5.5    81.3.4.4    Egress      LDP, BGP      POP-LDP
  4   299827  BGP      81.4.5.5    81.3.4.4    Egress      BGP           (null)

```

traceroute mpls bgp detail

user@host> traceroute mpls bgp fec detail

```

Probe options: retries 3, exp 7
Hop 2.2.1.81.rev.sfr.net (81.1.2.2) Depth 1
  Probe status: Success
  Parent: (null)
  Return code: Label switched at stack-depth 1
  Sender timestamp: 2013-03-22 05:55:19 PDT 822.99 msec
  Receiver timestamp: 2013-03-22 05:55:19 PDT 856.05 msec
  Response time: 33.06 msec
  MTU: Unknown
  Multipath type: IP bitmask
    Address Range 1: 127.0.0.64 ~ 127.0.0.127
  Label Stack:
    Label 1 Value 299824 Protocol LDP
    Label 2 Value 299276 Protocol BGP
  Fec-Stack-Sent: LDP, BGP
  Fec-Change:
    Operation: PUSH   Protocol RSVP

```


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 only 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.

Starting in Junos OS Release 14.1, the **transit l3vpn** statement is enabled by default on PTX Series Packet Transport Routers only.

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.

NOTE: The **ldp-p2mp** and **rsvp-p2mp** statements are not supported on MX series routers.

On an MX series router with redundant Routing Engines and enhanced-ip mode configuration, enabling the **ldp-p2mp** and **rsvp-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.

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.

NOTE: The **ldp-p2mp** and **rsvp-p2mp** statements are not supported on MX series routers.

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.

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.

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](#) | **2840**

RSVP Operational Commands

IN THIS CHAPTER

- `clear rsvp session` | 3356
- `clear rsvp statistics` | 3358
- `monitor label-switched-path` | 3360
- `ping mpls rsvp` | 3364
- `show rsvp interface` | 3371
- `show rsvp neighbor` | 3379
- `show rsvp route-session-id` | 3385
- `show rsvp pop-and-forward` | 3387
- `show rsvp session` | 3390
- `show rsvp session` | 3405
- `show rsvp statistics` | 3411
- `show rsvp version` | 3419
- `traceroute mpls rsvp` | 3423

clear rsvp session

List of Syntax

[Syntax on page 3356](#)

[Syntax \(EX and QFX Series Switches\) on page 3356](#)

Syntax

```
clear rsvp session
<all>
<connection-destination address>
<connection-source address>
<gracefully>
<logical-system (all | logical-system-name)>
<lsp-id identifier>
<name name>
<optimize-fast-reroute>
<tunnel-id identifier>
```

Syntax (EX and QFX Series Switches)

```
clear rsvp session
<connection-destination address>
<connection-source address>
<gracefully>
<lsp-id identifier>
<name name>
<optimize-fast-reroute>
<tunnel-id identifier>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Reset and restart Resource Reservation Protocol (RSVP) sessions.

Options

all—Clear all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.

connection-source *address*—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.

connection-destination *address*—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.

gracefully—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-id *identifier*—(Optional) LSP identifier (source port) for the RSVP sender template.

name *name*—(Optional) Reset and restart the specified RSVP session.

optimize-fast-reroute—(Optional) Begin fast reroute optimization.

tunnel-id *identifier*—(Optional) Tunnel identifier (destination port) for the RSVP session.

Required Privilege Level

clear

RELATED DOCUMENTATION

[clear mpls lsp | 3103](#)

[show rsvp session | 3390](#)

List of Sample Output

[clear rsvp session all on page 3357](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear rsvp session all
```

```
user@host> clear rsvp session all
```


clear rsvp statistics

List of Syntax

[Syntax on page 3358](#)

[Syntax \(EX Series Switches\) on page 3358](#)

Syntax

```
clear rsvp statistics  
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
clear rsvp statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Clear Resource Reservation Protocol (RSVP) packet and error statistics.

Options

none—Clear RSVP packet and error statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

RELATED DOCUMENTATION

[show rsvp statistics](#) | [3411](#)

List of Sample Output

[clear rsvp statistics on page 3359](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear rsvp statistics
```

```
user@host> clear rsvp statistics
```


monitor label-switched-path

Syntax

```
monitor label-switched-path lsp-name
<logical-system (logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Logical system support introduced in Junos OS Release 9.4.

Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Display the real-time status of the specified RSVP label-switched path (LSP). You can also use this command to monitor LSPs configured within logical systems.

Options

logical-system (*logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-name—Name of the LSP.

Additional Information

You can track the amount of traffic traversing an RSVP LSP and observe its essential parameters, such as uptime, ingress and egress addresses, labels, routes, and ports. Values are typically sampled every second. The display also allows you to scroll to other currently running LSPs. You cannot use this command to display information about static LSPs or LDP-signaled LSPs.

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the **c** key. To control the output of the **monitor label-switched-path** command while it is running, use the keys listed in [Table 120 on page 3360](#). The keys are not case-sensitive.

Table 120: 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.

Table 120: Output Control Keys for the monitor label-switched-path Command (*continued*)

Key	Action
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 3362](#)**Output Fields**

[Table 121 on page 3361](#) describes the output fields for the **monitor label-switched-path** command. Output fields are listed in the approximate order in which they appear.

Table 121: monitor label-switched-path Output Fields

Field Name	Field Description
(1)	<p>Displays the following information:</p> <ul style="list-style-type: none"> • hostname—Name of the router. • Seconds—Time elapsed since this display was started. • Time—Current local time.
(2)	<p>Delay—Length of the time delay, in milliseconds, required to obtain the information in the monitor display. The first number shows the current sampling delay. The second number shows the shortest delay recorded to date. The third number shows the worst delay recorded to date. This delay can vary substantially depending on the system load.</p>
(3)	<p>Displays the following:</p> <ul style="list-style-type: none"> • To—Destination address of the LSP. • From—Originating address of the LSP. • State—Current state of the LSP: Up or Down.
(4)	<p>Displays the following:</p> <ul style="list-style-type: none"> • LSPName—Name of the LSP. • Type—Type of LSP: Ingress, Egress, or Transit.

(12)

(13)Next='n', Prev='p', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c',
LSP='l'

ping mpls rsvp

Syntax

```
ping mpls rsvp
<lsp-name>
<count count>
<destination address>
<detail>
<dynamic-bypass>
<egress egress-address>
<exp forwarding-class>
<interface interface-name>
<logical-system (all | logical-system-name)>
<manual-bypass>
<multipoint>
<size bytes>
<source source-address>
<standby standby-path-name>
<sweep>
```

Release Information

Command introduced before Junos OS Release 7.4.

The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.

The **size** and **sweep** options were introduced in Junos OS Release 9.6.

The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.

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.

NOTE: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

List of Sample Output

[ping mpls rsvp \(Echo Reply Received\) on page 3366](#)
[ping mpls rsvp \(Echo Reply with Error Code\) on page 3367](#)
[ping mpls rsvp detail on page 3367](#)
[ping mpls rsvp multipoint egress detail count on page 3367](#)
[ping mpls rsvp multipoint detail count on page 3367](#)
[ping mpls rsvp destination detail count size on page 3368](#)
[ping mpls rsvp destination detail sweep size on page 3369](#)

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
```


show rsvp interface

List of Syntax

[Syntax on page 3371](#)

[Syntax \(EX Series Switches\) on page 3371](#)

Syntax

```
show rsvp interface
<brief | detail | extensive>
<instance instance-name>
<link-management>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp interface
<brief | detail | extensive>
<link-management>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance option added in Junos OS Release 15.1 for the MX Series.

Description

Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics. The RSVP input/input module collects statistics for certain events on a per-interface basis. Most of these events were tracked on a routing-instance basis in Junos OS releases earlier than Release 17.2. The **show rsvp interface detail** command displays these event counters under the **Events** section of the output only when the values of these fields are higher than zero. These statistics are also maintained at the global level (per routing-instance) and are also displayed in the output of the **show rsvp statistics** command.

Options

none—Display standard information about the status of RSVP-enabled interfaces and packet statistics.

brief | detail | extensive | link-management—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display RSVP status information for the specified instance. If **instance-name** is omitted, RSVP status information is displayed for the master instance.

link-management—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

[show rsvp interface brief on page 3375](#)

[show rsvp interface detail on page 3375](#)

[show rsvp interface extensive on page 3376](#)

[show rsvp interface link-management on page 3377](#)

[show rsvp interface detail RSVP interface: 9 active on page 3377](#)

Output Fields

[Table 122 on page 3372](#) lists the output fields for the **show rsvp interface** command. Output fields are listed in the approximate order in which they appear.

Table 122: show rsvp interface Output Fields

Field Name	Field Description	Level of Output
RSVP interface	Number of interfaces on which RSVP is active. Each interface has one line of output.	All levels
Interface	Name of the interface.	All levels
Index	Index of the interface.	detail
State	State of the interface. <ul style="list-style-type: none"> • Disabled—No traffic engineering information is displayed. • Down—Interface is not operational. • Enabled—Displays traffic engineering information. • Up—Interface is operational. 	All levels
NoAuthentication	Interface does not support RSVP authentication.	detail
NoAggregate	Interface does not support refresh reduction.	detail
NoReliable	Interface does not support refresh reduction message ID extension.	detail
NoLinkProtection	Interface does not support link protection.	detail

Table 122: show rsvp interface Output Fields (continued)

Field Name	Field Description	Level of Output
HelloInterval	Frequency at which RSVP hellos are sent on this interface (in seconds). Prior to Junos OS Release 18.2R2, when the no-interface-hello statement is configured under the [edit protocols rsvp] hierarchy, and there is no interface-specific configuration for the hello interval, the HelloInterval output field displayed the default hello interval time of 9 seconds. Starting in Junos OS Release 18.2R2, with a similar configuration, the HelloInterval output field displays 0 as the hello interval.	detail
Address	IP address of the local interface.	detail
Active control channel	Next-hop link address to transmit messages.	None specified
TElink	Traffic-engineered links that are managed by the peer they are associated with.	None specified
Active resv	Number of reservations that are actively reserving bandwidth on the interface.	All levels
PreemptionCnt	Number of times an RSVP session was preempted on this interface.	detail
Update threshold	Percentage change in reserved bandwidth to trigger an IGP update.	detail
Subscription	User-configured subscription factor.	All levels
Actual	Available RSVP bandwidth that is recalculated after considering SPRING bandwidth utilization.	extensive
bc number	Bandwidth allocated for the specified bandwidth constraint.	extensive
ct number	Bandwidth allocated for the specified class type.	extensive
Static BW	Total interface bandwidth, in bps.	All levels
Available BW	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	al levels
Reserved BW	Currently reserved bandwidth, in bps.	All levels
SoftPreemptionCnt	Number of times a soft preemption occurred on this interface. This number is not included in the PreemptionCnt value.	detail

Table 122: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Overbooked BW	Currently overbooked bandwidth, in bps, by class type (ct0 through ct3).	detail
Highwater mark	Highest bandwidth that has ever been reserved on this interface, in bps.	brief
PacketType	Type of RSVP packet.	detail
Total Sent	Total number of packets sent.	detail
Total Received	Total number of packets received since RSVP was enabled.	detail
Last 5 seconds Sent	Number of packets sent in the last 5 seconds.	detail
Last 5 seconds Received	Number of packets received in the last 5 seconds.	detail
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	detail
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	detail
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	detail
Resv	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	detail
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	detail
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.	detail
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	detail
Ack	Acknowledge message for refresh reductions.	detail

Table 122: show rsvp interface Output Fields (continued)

Field Name	Field Description	Level of Output
Srefresh	Summary refresh messages.	detail
EndtoEnd RSVP	Statistics for the number of end-to-end RSVP messages sent.	detail
Queue	CoS transmit queue number and its associated forwarding class designation.	extensive
TxRate	Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.	extensive
Priority	Weight of the queue relative to other configured queues, in percentage.	extensive
queue-priority-value	Low, High, None, or Exact. None indicates no rate limiting. Exact indicates the queue transmits at the configured rate only.	extensive

Sample Output

show rsvp interface brief

```
user@host> show rsvp interface brief
```

```

RSVP interface: 1 active
           Active Subscr- Static   Available   Reserved   Highwater
Interface  State resv   iption  BW        BW         BW         mark
de0.0      Up        1      23%    10Mbps    989.992kbps 1.31Mbps   1.31Mbps

```

show rsvp interface detail

Starting in Junos OS Release 15.2, this command also shows conditional PathTear statistics and Node Hellos.

```
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, SoftPreemptionCnt 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
Node Hello          100          100        0           0
PathTear(Condl.)    0           3         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, SoftPreemptionCnt 0, Update threshold 10%
Subscription 100%, Actual 60%
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 interface detail RSVP interface: 9 active

user@host> show rsvp interface detail RSVP interface: 9 active

```

fxp0.0 Index 4, State Dis/Up
NoAuthentication, Aggregate, Reliable, NoLinkProtection HelloInterval 9(second)
Address 10.9.148.47
Event                               Count
bad packet length                   1
bad packet version                   1
authentication fail                  1
bad checksum                         1

```


bad packet format	1
rcv pkt disabled intf	1
state timeout	1
message out-of-order	1
unknown ack	1
unknown nack	1
received nack	1
send failure	1

PacketType		Total	Last 5 seconds			
Sent	Received		Sent	Received		
Path		0	0	0	0	0
PathErr		0	0	0	0	0
PathTear		0	0	0	0	0
Resv		0	0	0	0	0
ResvErr		0	0	0	0	0
ResvTear		0	0	0	0	0
ResvConf		0	0	0	0	0
Bundle		0	0	0	0	0
Hello		0	0	0	0	0
Ack		0	0	0	0	0
Srefresh		0	0	0	0	0
Notify		0	0	0	0	0
Unknown		0	0	0	0	0
EndtoEnd RSVP		0	0	0	0	0
Backup Path		0	0	0	0	0
Cnd PathTear		0	0	0	0	0

show rsvp neighbor

List of Syntax

[Syntax on page 3379](#)

[Syntax \(EX Series Switches\) on page 3379](#)

Syntax

```
show rsvp neighbor  
<brief | detail | extensive>  
<instance instance-name>  
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp neighbor  
<brief | detail>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance option added in Junos OS Release 15.1 for the MX Series.

Description

Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.

Options

none—Display standard information about RSVP neighbors.

brief | detail—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display the RSVP neighbor information for the specified instance. If **instance-name** is omitted, RSVP neighbor 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.

Required Privilege Level

view

List of Sample Output

[show rsvp neighbor on page 3383](#)

[show rsvp neighbor detail on page 3384](#)

Output Fields

Table 123 on page 3380 lists the output fields for the **show rsvp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 123: show rsvp neighbor Output Fields

Field Name	Field Description	Level of Output
RSVP neighbor	Number of neighbors that the routing device has learned of. Each neighbor has one line of output.	All levels
via	Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected.	detail
Address	Address of a learned neighbor.	All levels
Idle	Length of time the neighbor has been idle, in seconds. NOTE: Until Junos OS Release 15.1, in the output of the show rsvp neighbor command, the value under the Idle field immediately reflects the changed idle time when a link in the neighboring router is brought down. Starting with Junos OS Release 15.2, a router does not declare a neighbor as idle when a hello adjacency exists and has not timed out. When an interface is brought down, RSVP brings down the neighbor because of the notification it receives from IGP. The reason for considering the IGP-down notification is to support BFD-triggered fast reroute (FRR) and RSVP-TE is not directly a client for BFD notifications. When RSVP brings down the neighbor, the input/output process is not impacted. As a result, the idle time in the output of the show command is not immediately updated.	All levels
Up/Dn	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	All levels
Up cnt and Down cnt	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	detail

Table 123: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
status	<p>State of the RSVP neighbor:</p> <ul style="list-style-type: none"> • Up—Routing device can detect RSVP Hello messages from the neighbor. • Down—Routing device has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP Hello messages sent by the neighbor. • Restarting—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled. • Restarted—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures. • Dead—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down. 	detail
LastChange	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is hh:mm:ss .	All levels
Last changed time	Time elapsed since the neighbor state changed either from up to down or from down to up.	detail
HelloInt	Frequency at which RSVP hellos are sent on this interface (in seconds).	All levels
HelloTx/Rx	Number of hello packets sent to and received from the neighbor.	All levels
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	detail
Message received	Number of Path and Resv messages that this routing device has received from the neighbor.	detail
Remote Instance	Identification provided by the remote routing device during Hello message exchange.	detail
Local Instance	Identification sent to the remote routing device during Hello message exchange.	detail

Table 123: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Refresh reduction	<p>Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. Refresh reduction can have the following values:</p> <ul style="list-style-type: none"> • operational—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961. • incomplete—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices. • not operational—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions. 	detail
Remote end	<p>Neighboring routing device's status with regard to refresh reduction:</p> <ul style="list-style-type: none"> • enabled—Remote routing device has requested refresh reduction during RSVP message exchanges. • disabled—Remote routing device does not require refresh reduction. 	detail
Pop label	Pop labels of the RSVP-TE pop-and-forward LSP tunnels.	detail
Ack-extension	<p>An RSVP refresh reduction extension:</p> <ul style="list-style-type: none"> • enabled—Both local and remote routing devices support the ack-extension (RFC 2961). • disabled—Remote routing device does not support the ack-extension. 	detail
Link protection	<p>Status of the MPLS fast reroute mechanism that protects traffic from link failure:</p> <ul style="list-style-type: none"> • enabled—Link protection feature has been turned on, protecting the neighbor with a bypass LSP. • disabled—No link protection feature has been enabled for this neighbor. 	detail
LSP name	Name of the bypass LSP.	detail

Table 123: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Bypass LSP	Status of the bypass LSP. It can have the following values: <ul style="list-style-type: none"> • does not exist—Bypass LSP is not available. • connecting—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment. • operational—Bypass LSP is up and running. • down—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path. 	detail
Backup routes	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	detail
Backup LSPs	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	detail
Bypass explicit route	Explicit route object's (ERO) path that is taken by the bypass LSP.	detail
Restart time	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	detail
Recovery time	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	detail

Sample Output

show rsvp neighbor

user@host> **show rsvp neighbor**

```

RSVP neighbor: 2 learned
Address           Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203    0  3/2    13:01      3    366/349
192.168.207.207    0  1/0    22:49      3    448/448

```


show rsvp neighbor detail

Starting in Junos OS Release 16.1, this command also shows whether enhanced FRR procurers are enabled on the neighbor. Neighbors with Point of Local Repair (PLR) or Node Protecting Merge Point (NP-MP) also show the Hellos sent /received count.

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
  Enhanced FRR local protection: enabled
    LSPs (total 76): Phop 0, PPhop 0, Nhop 76, NNhop 0
  Pop Label: 299808(unprotected) 299840(link-protected)
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 route-session-id

Syntax

```
show rsvp route-session-id
```

Release Information

Command introduced in Junos OS Release 16.1 for the MX Series.

Description

Display the session ID and the version information associated with the ingress route added by the Resource Reservation Protocol (RSVP) in the inet.3 table.

Session ID is a pre-populated identifier used for indirect next hops in BGP Prefix Independent Convergence (PIC) enabled router. Session ID is used to identify the session or path.

NOTE: `protect core` configuration is not required to display the `route-session-id`.

Options

none—Validate and display RSVP route session details.

Required Privilege Level

view

List of Sample Output

[show rsvp route-session-id on page 3386](#)

Output Fields

[Table 124 on page 3385](#) describes the output fields for the `show rsvp route-session-id` command. Output fields are listed in the approximate order in which they appear.

Table 124: show rsvp route-session-id Output Fields

Field Name	Field Description
Ingress Route Destination	Destination (egress routing device) of the session.
Ingress Route Preference	RSVP preference value of the ingress session.
Ingress Route Metric 1	Metric 1 associated with the RSVP ingress route.

Table 124: show rsvp route-session-id Output Fields (*continued*)

Field Name	Field Description
Ingress Route Metric 2	Metric 2 associated with the RSVP ingress route.
Ingress Route Session ID	Session ID associated with the RSVP ingress route.
Version	Version number associated with the RSVP ingress route.

Sample Output

show rsvp route-session-id

user@host> **show rsvp route-session-id**

```
RSVP Ingress Route Session ID Database:
=====
Ingress Route Destination: 1.1.1.5/32
Ingress Route Preference: 7
Ingress Route Metric 1:    20,          Metric 2: 0
Ingress Route Session ID: 0x146,       Version: 0
```


show rsvp pop-and-forward

Syntax

```
show rsvp pop-and-forward
<brief | detail | extensive>
<instance routing-instance-name>
<label label>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 18.1R1 on MX Series routers, PTX Series routers, and vMX series routers.

Description

Display RSVP-TE pop-and-forward LSP tunnel information. This information includes the set of in-labels (one-hop pop label or a delegation label), the number of session using each label and the next segment-label (if there is another delegation hop downstream), and whether the in-label is used for unprotected or protected LSPs.

Options

none—Display the standard level of information for the RSVP-TE pop-and-forward LSP tunnels.

brief | detail | extensive—(Optional) Display the desired level of output. The brief option is the default level of output.

The detail option provides more information about the hops in a delegation segment (whether its one-hop or multi-hop).

The extensive option lists the set of LSPs that are using a given pop or delegation label.

instance *routing-instance-name*—(Optional) Display the RSVP-TE pop-and-forward LSP tunnel information for the specified routing instance.

label *label*—(Optional) Display the RSVP-TE pop-and-forward LSP tunnel information for the specified label.

logical-system (all | *logical-system-name*)—(Optional) Display the RSVP-TE pop-and-forward LSP tunnel information for all or the specified logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

List of Sample Output

[show rsvp pop-and-forward on page 3388](#)

[show rsvp pop-and-forward extensive on page 3388](#)

[show rsvp pop-and-forward label on page 3389](#)

Sample Output

show rsvp pop-and-forward

user@host> show rsvp pop-and-forward

RSVP pop-and-forward: 2 shared labels				
Label-in	Hop-count	Next-segment-label	Protection	Session-count
299840	3	299808	unprotected	100
299872	3	299824	unprotected	50

show rsvp pop-and-forward extensive

user@host> show rsvp pop-and-forward extensive

RSVP pop-and-forward: 2 shared labels				
299840 (shared-label)				
Next-segment-label: 299808, Hop-count: 3				
Protection: unprotected, Session-count: 2				
Segment-id:				
Hop 1: 70.1.1.2(label=299808)				
Hop 2: 92.1.1.1(label=299808)				
Hop 3: 93.1.1.2				
Segment route:				
Primary: 70.1.1.2, OutIf: ge-0/0/2.0				
Lsp-session list (name, dest-ip, sender-ip, lsp-id):				
pop1, 10.10.10.10, 2.2.2.2, 2				
pop2, 10.10.10.10, 2.2.2.2, 1				
299872 (shared-label)				
Next-segment-label: 299824, Hop-count: 3				
Protection: unprotected, Session-count: 4				


```
Segment-id:
  Hop 1: 70.1.1.2(label=299808)
  Hop 2: 92.1.1.1(label=299808)
  Hop 3: 93.1.1.2
Segment route:
  Primary: 70.1.1.2, OutIf: ge-0/0/2.0
Lsp-session list (name, dest-ip, sender-ip, lsp-id):
  pop147, 9.9.9.9, 2.2.2.2, 1
  pop148, 9.9.9.9, 2.2.2.2, 1
  pop150, 9.9.9.9, 2.2.2.2, 1
  pop149, 9.9.9.9, 2.2.2.2, 1
```

show rsvp pop-and-forward label

user@host> **show rsvp pop-and-forward label 299872**

RSVP pop-and-forward: 2 shared labels				
Label-in	Hop-count	Next-segment-label	Protection	Session-count
299872	3	299824	unprotected	4

show rsvp session

List of Syntax

[Syntax on page 3390](#)

[Syntax \(EX and QFX Series Switches\) on page 3390](#)

Syntax

```
show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<externally-provisioned>
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
<lsp-type>
<name session-name>
<p2mp>
<session-type>
<statistics>
<te-link te-link>
```

Syntax (EX and QFX Series Switches)

```
show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<externally-provisioned>
<interface interface-name>
<lsp-type>
<name session-name>
<p2mp>
<session-type>
<statistics>
<te-link te-link>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

externally-provisioned option added in Junos OS Release 13.3.

Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.

instance option added in Junos OS Release 15.1 for the MX Series.

Description

Display information about RSVP sessions.

Options

none—Display standard information about all RSVP sessions.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

bidirectional | unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.

bypass—(Optional) Display RSVP sessions for bypass LSPs.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

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.

RSVP reserves resources only for outgoing LSPs of an interface. Because resources are not reserved for incoming LSPs, the **show RSVP sessions interface *interface-name*** command output displays only those RSVP sessions whose next hops correspond to the specified interface.

To identify the number of RSVP sessions formed over the uplink interface on the egress label-switching router (LSR), you can use the following command:

```
user@host> show RSVP session egress extensive | match "PATH rcvfrom:" | match interface-name | count
```

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.

To identify the number of RSVP sessions formed over the uplink interface on the egress label-switching router (LSR), you can use the following command:

```
user@host> show rsvp session egress extensive | match "PATH rcvfrom:" | match interface-name
| count
```

- **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](#) | [3356](#)

List of Sample Output

[show rsvp session on page 3397](#)

[show rsvp session statistics on page 3398](#)

[show rsvp session detail on page 3398](#)

[show rsvp session detail \(When Egress Protection Is in Standby Mode\) on page 3399](#)

[show rsvp session detail \(When Egress Protection Is in Effect During a Local Repair\) on page 3399](#)

[show rsvp session detail \(Path MTU Output Field\) on page 3400](#)

[show rsvp session detail \(GMPLS\) on page 3400](#)

[show rsvp session extensive on page 3401](#)

[show rsvp session extensive transit on page 3402](#)

[show rsvp session p2mp \(Ingress Router\) on page 3403](#)

[show rsvp session p2mp \(Transit Router\) on page 3403](#)

Output Fields

[Table 125 on page 3393](#) describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 125: show rsvp session Output Fields

Field Name	Field Description	Level of Output
Ingress RSVP	Information about ingress RSVP sessions.	detail
Ingress RSVP	Information about ingress RSVP sessions. Each session has one line of output.	All levels
Egress RSVP	Information about egress RSVP sessions.	All levels
Transit RSVP	Information about the transit RSVP sessions.	All levels
P2MP name	(Appears only when the p2mp option is specified). Name of the point-to-multipoint LSP path.	All levels
P2MP branch count	(Appears only when the p2mp option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels
To	Destination (egress routing device) of the session.	All levels
From	Source (ingress routing device) of the session.	All levels
State	State of the path: Up , Down , or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress routing device) of the LSP.	detail
From	Source (ingress routing device) of the session.	detail
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either Up , Dn (down), or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	brief detail
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
Active Route	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail
LSPname	Name of the LSP.	brief detail

Table 125: show rsvp session Output Fields (continued)

Field Name	Field Description	Level of Output
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices. LSPpath can also indicate when a graceful LSP deletion has been triggered.	detail
Bypass	(Egress routing device) Destination address for the bypass LSP.	detail
Bidir	(When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.	detail
Bidirectional	(When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.	detail
Upstream label in	(When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.	detail
Upstream label out	(When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.	detail
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	detail
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	detail
Resv style or Style	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail
Label in	Incoming label for this LSP.	brief detail
Label out	Outgoing label for this LSP.	brief detail
Time left	Number of seconds remaining in the lifetime of the reservation.	brief detail

Table 125: show rsvp session Output Fields (continued)

Field Name	Field Description	Level of Output
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
DiffServ info	Indicates whether the LSP is a multiclass LSP (multiclass diffServ-TE LSP) or a Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP).	detail
bandwidth	Bandwidth for each class type (ct0 , ct1 , ct2 , or ct3).	detail
Port number	Protocol ID and sender/receiver port used in this RSVP session.	detail
Attrib flags	Non-PHP indicates that ultimate hop popping has been requested by the LSP using this RSVP session	extensive
FastReroute desired	Fast reroute has been requested by the ingress routing device.	detail
Soft preemption desired	Soft preemption has been requested by the ingress routing device.	detail
FastReroute desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.	detail extensive
Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.	detail extensive
Node/Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	detail extensive

Table 125: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type	<p>LSP type:</p> <ul style="list-style-type: none"> • Link protected LSP—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Node/Link protected LSP—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Protection down—LSP is not currently protected. • Bypass LSP—LSP that is used to protected one or more user LSPs in case of link failure. • Backup LSP at Point-of-Local-Repair (PLR)—LSP that has been temporarily established to protected a user LSP at the ingress of a failed link. • Backup LSP at Merge Point (MP)—LSP that has been temporarily established to protected a user LSP at the egress of a failed link. 	detail extensive
New bypass	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	extensive
Link protection up, using <i>bypass-name</i>	Link protection (the bypass name is also displayed) has been activated for the LSP.	extensive
Creating backup LSP, link down	A link down event occurred, and traffic is being switched over to the bypass LSP.	extensive
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	extensive
Path mtu	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the allow-fragmentation statement configured at the [edit protocols mpls path-mtu] hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	detail

Table 125: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Egress protection PLR as protector	RSVP state on the Protector or the point-of-local-repair (PLR) routing device: <ul style="list-style-type: none"> • Active— Egress protection is available at the Protector or the PLR routing device. • In Use— Local repair has been completed. 	detail
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.	detail
Adspec	MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.	detail
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.	detail
Explct route	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail
Record route	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail

Sample Output

show rsvp session

user@host> **show rsvp session**

```
Ingress RSVP: 1 sessions
To          From          State  Rt  Style  Labelin  Labelout  LSPname
10.255.245.214 10.255.245.212 AdminDn 0  1  FF      -      22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To          From          State  Rt  Style  Labelin  Labelout  LSPname
```



```

10.255.245.194 10.255.245.195 Up      0  1 FF  39811      - Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up      0  1 FF      3      - pro3-ba
Total 2 displayed, Up 2, Down 0

```

Transit RSVP: 1 sessions

```

To           From           State Rt Style Labelin Labelout LSPname
10.255.245.198 10.255.245.197 Up      0  1 SE  100000      3 pro3-de
Total 1 displayed, Up 1, Down 0

```

show rsvp session statistics

user@host> show rsvp session statistics

```

Ingress RSVP: 2 sessions
To           From           State  Packets    Bytes    LSPname
10.255.245.24 10.255.245.22  Up      0          0      pro3-bd
10.255.245.24 10.255.245.22  Up    44868    2333136  pro3-bd-2
Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To           From           State  Packets    Bytes    LSPname
10.255.245.22 10.255.245.24  Up      0          0      pro3-db
10.255.245.22 10.255.245.24  Up      0          0    pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show rsvp session detail

user@host> show rsvp session detail

```

Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ctl 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: <ctl 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: <ctl 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-rl-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

Starting in Junos OS Release 16.1, this command includes additional details for both the incoming and outgoing Path and Resv messages. The information includes the internal message handle and revision number, as well as the message ID included by the neighbor in the signaling message.

user@host> **show rsvp session extensive**

```

Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 1 sessions

16.0.0.5
  From: 16.0.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lto5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299856
  Resv style: 1 SE, Label in: 299776, Label out: 299856
  Time left: 123, Since: Sat Nov 29 10:39:15 2014
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 52631 protocol 0
  PATH rcvfrom: 16.1.2.1 (ge-0/0/0.0) 2 pkts
  incoming message handle: P-1/2, ID: 0xc82fd7/322
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 16.2.4.4 (ge-0/0/1.0) 1 pkts
  outgoing message state: refreshing, ID: 0xcacec0/22
  RESV rcvfrom: 16.2.4.4 (ge-0/0/1.0) 1 pkts, Entropy label: Yes
  incoming message handle: R-2/1, ID: 0xc82f3e/217
  RESV
  outgoing message state: refreshing, ID: 0xcacec0/17
  Explct route: 16.2.4.4 16.99.0.5

```



```
Record route: 16.1.2.1 <self> 16.2.4.4 16.99.0.5
Total 1 displayed, Up 1, Down 0
```

show rsvp session extensive transit

Starting in Junos OS Release 16.1, this command also shows node-related details, including whether enhanced local protection is enabled for the LSP and whether the node is a merge point. If the latter is true, both the IP address of the Point of Local Repair (PLR) and the status (LP-MP, NP-MP, or Non-MP) are shown.

user@host> **show rsvp session extensive transit**

```
From: 81.1.1.1, LSPstate: Up, ActiveRoute: 0
LSPname: A-D-1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 299776
Resv style: 1 SE, Label in: 299776, Label out: 299776
Time left: 117, Since: Tue May 6 08:39:44 2014
Tspec: rate 700Mbps size 700Mbps peak Infbps m 20 M 1500
Port number: sender 1 receiver 24131 protocol 0
Node/Link protection desired
Type: Node/Link protected LSP, using Bypass->81.2.3.3->81.3.4.4
    2 May 6 08:39:47 Node protection up, using Bypass->81.2.3.3->81.3.4.4
    1 May 6 08:39:44 New bypass Bypass->81.2.3.3->81.3.4.4
Enhanced Local Protection: Enabled, LP-MP for 81.2.2.2, NP-MP for 81.1.1.1
PATH rcvfrom: 81.1.2.1 (lt-0/2/0.201) 5371 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 81.2.3.3 (lt-0/2/0.203) 5374 pkts
RESV rcvfrom: 81.2.3.3 (lt-0/2/0.203) 5372 pkts, Entropy label: No
Record route: 81.1.2.1 <self> 81.3.3.3 (node-id) 81.2.3.3 81.4.4.4 (node-id)
81.3.4.4
Total 1 displayed, Up 1, Down 0
```

If enhanced FRR is not enabled (either because it is disabled on the router itself or one of the neighbors along the LSP path does not support it), either of the following lines might be displayed:

```
Enhanced Local Protection: Disabled, Reason: User Config
```

```
Enhanced Local Protection: Disabled, Reason: Backward Compatibility
```

If enhanced FRR is not enabled and the router is not an MP, the following line is displayed:

Enhanced Local Protection: Enabled, Non-MP

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 session

Syntax

```
show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<down | up>
<interface interface-name>
<lsp-type>
<name session-name>
<session-type>
<statistics>
<te-link te-link>
```

Release Information

Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display information about Resource Reservation Protocol (RSVP) sessions.

Options

none—Display standard information about all RSVP sessions.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

bidirectional | unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

interface *interface-name*—(Optional) Display RSVP sessions for the specified interface only.

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.

session-type—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this switch.
- **ingress**—Sessions that originate from this switch.

- **transit**—Sessions that transit through this switch.

statistics—(Optional) Display packet statistics.

te-link *te-link*—(Optional) Display sessions with reservations on the specified traffic-engineered link name.

Required Privilege Level

view

RELATED DOCUMENTATION

[Example: Configuring MPLS on EX8200 and EX4500 Switches | 59](#)

[Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect \(CLI Procedure\) | 95](#)

[Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) | 89](#)

[Configuring MPLS on EX8200 and EX4500 Provider Switches \(CLI Procedure\) | 99](#)

List of Sample Output

[show rsvp session on page 3408](#)

[show rsvp session statistics on page 3409](#)

[show rsvp session detail on page 3409](#)

[show rsvp session extensive on page 3410](#)

Output Fields

[Table 126 on page 3406](#) describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 126: show rsvp session Output Fields

Field Name	Field Description	Level of Output
Ingress RSVP	Information about ingress RSVP sessions.	detail
Ingress RSVP	Information about ingress RSVP sessions. Each session has one line of output.	All levels
Egress RSVP	Information about egress RSVP sessions.	All levels
Transit RSVP	Information about the transit RSVP sessions.	All levels
To	Destination (egress switch) of the session.	All levels
From	Source (ingress switch) of the session.	All levels

Table 126: show rsvp session Output Fields (continued)

Field Name	Field Description	Level of Output
State	State of the path: Up , Down , or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress switch) of the LSP.	detail
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either Up , Dn (down), or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	brief, detail
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
ActiveRoute	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail
LSPname	Name of the LSP.	brief, detail
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit switches. LSPpath can also indicate when a graceful LSP deletion has been triggered.	detail
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	detail
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	detail
Resv style or Style	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail

Table 126: show rsvp session Output Fields (continued)

Field Name	Field Description	Level of Output
Label in	Incoming label for this LSP.	brief, detail
Label out	Outgoing label for this LSP.	brief, detail
Time left	Number of seconds remaining in the lifetime of the reservation.	brief, detail
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
Port number	Protocol ID and sender/receiver port used in this RSVP session.	detail
Creating backup LSP, link down	A link down event occurred, and traffic is being switched over to the bypass LSP.	extensive
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	extensive
PATH rcvfrom	Address of the previous-hop (upstream) switch or client, interface the neighbor used to reach this switch, and number of packets received from the upstream neighbor.	detail

Sample Output

show rsvp session

```
user@switch> 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@switch> 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@switch> 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: <ctl 300kbps>
```

```
Port number: sender 1 receiver 15876 protocol 0
```

```
PATH rcvfrom: localclient
```

```
Adspec: sent MTU 1500
```

```
PATH sentto: 192.168.37.16 (tl-0/2/1.0) 1 pkt
```


show rsvp session extensive

user@switch> **show rsvp session extensive**

```
8.8.8.8
  From: 9.9.9.9, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp_to_240, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 322832
  Resv style: 1 FF, Label in: -, Label out: 322832
  Time left:    -, Since: Thu Feb 26 16:25:39 2009
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 44542 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 3.3.3.2 (xe-0/1/0.0) 238 pkts
  RESV rcvfrom: 3.3.3.2 (xe-0/1/0.0) 234 pkts
  Explct route: 3.3.3.2 4.4.4.2
```


show rsvp statistics

List of Syntax

[Syntax on page 3411](#)

[Syntax \(EX Series Switches\) on page 3411](#)

Syntax

```
show rsvp statistics  
<instance instance-name>  
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp statistics
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

instance option added in Junos OS Release 15.1 for the MX Series.

Description

Display Resource Reservation Protocol (RSVP) packet and error statistics. The RSVP input/output module collects statistics for certain events on a per-interface basis. Most of these events were tracked on a routing-instance basis in Junos OS releases earlier than Release 17.2. The "show rsvp interface detail" command displays these event counters under the Events section of the output only when the values of these fields are higher than zero. These statistics are also maintained at the global level (per routing-instance) and are also displayed in the output of the "show rsvp statistics" command.

Options

none—Display RSVP packet and error statistics.

instance *instance-name*—(Optional) Display RSVP packet and error statistics for the specified instance. If **instance-name** is omitted, RSVP statistics are 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.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear rsvp statistics](#) | [3358](#)

List of Sample Output

[show rsvp statistics on page 3415](#)

[show rsvp statistics on page 3416](#)

Output Fields

[Table 127 on page 3412](#) describes the output fields for the **show rsvp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 127: show rsvp statistics Output Fields

Field Name	Field Description
Packet Type	Statistics about different RSVP messages.
Total Sent	Total number of packets sent since RSVP was enabled.
Total Received	Total number of packets received since RSVP was enabled.
Last 5 seconds Sent	Total number of packets sent in the last 5 seconds.
Last 5 seconds Received	Number of packets received in the last 5 seconds.
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.
Resv FF	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
Resv WF	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
Res SE	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.

Table 127: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.
ResvConf	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
Ack	Acknowledge message for refresh reductions.
SRefresh	Summary refresh messages.
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.
EndtoEnd RSVP	Statistics for the number of End-to-end RSVP messages.
Errors	Statistics about errored RSVP packets.
Rcv pkt bad length	The packet was not processed because its length is inappropriate.
Rcv pkt unknown type	The packet is not one of the well-known RSVP types, as defined in RFC 2205, <i>Resource ReSerVation Protocol (RSVP)</i> .
Rcv pkt bad version	The packet is not an RSVP version 1 packet.
Rcv pkt auth fail	The packet failed authentication checks.
Rcv pkt bad checksum	The RSVP checksum check failed.
Rcv pkt bad format	General packet processing failed because the packet was badly formed.
Memory allocation fail	An internal resource failure occurred.
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.

Table 127: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
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 recv 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.
Err msg loop detected	Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object.

Sample Output

show rsvp statistics

Starting in Junos OS Release 16.1, this command also shows conditional PathTear statistics and the number of times an LSP state has been retained because of Link Protecting Merge Point (LP-MP) or Node Protecting Merge Point (NP-MP) determination.

```
user@host> show rsvp statistics
```

PacketType	Total		Last 5 seconds	
	Sent	Received	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
Bundle	455	378	0	0
Ack	682	1414	0	0
SRefresh	395198	236030	5	2
Hello	578809	578221	4	4
EndtoEnd RSVP	0	0	0	0
Node Hello	50	50	0	0
PathTear(Cond1.)	0	3	0	0
Errors	Total		Last 5 seconds	
Rcv pkt bad length	0		0	
Rcv pkt unknown type	0		0	
Rcv pkt bad version	0		0	
Rcv pkt auth fail	0		0	
Rcv pkt bad checksum	0		0	
Rcv pkt bad format	0		0	
Memory allocation fail	0		0	
No path information	10		0	
Resv style conflict	0		0	
Port conflict	0		0	
Resv no interface	0		0	
PathErr to client	38		0	
ResvErr to client	0		0	
Path timeout	8		0	
Resv timeout	57		0	

Message out-of-order	0	0
Unknown ack msg	2978	0
Recv nack	86	0
Recv duplicated msg-id	5	0
No TE-link to recv Hop	0	0
Rcv pkt disabled interface	0	0
Transmit buffer full	0	0
Transmit failure	0	0
Receive failure	0	0
P2MP RESV discarded by appl	0	0
Rate limit	306	0
Err msg loop detected	0	0
MP Path LP-avail rcved	0	0
MP Path NP-avail rcved	0	0
PLR bk RSB life ext	0	0
MP bk PSB life ext	0	0
LP-MP state retained on failure	0	0
NP-MP state retained on failure	0	0
Fast refresh skipped	0	0
MP bk Srefresh Nack rcved	0	0
RSB life extended for nh FRR	0	0

show rsvp statistics

user@host> show rsvp statistics

PacketType	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Path	21	0	0	0
PathErr	0	4	0	0
PathTear	9	0	0	0
Resv	0	9	0	0
ResvErr	0	0	0	0
ResvTear	0	2	0	0
ResvConf	0	0	0	0
Bundle	28	2	0	0
Hello	172814	172802	5	5
Ack	11	12	0	0
Srefresh	142	143	0	0
Notify	0	0	0	0
Unknown	0	0	0	0
EndtoEnd RSVP	0	0	0	0
Backup Path	0	0	0	0
Backup Tear	0	0	0	0

Cnd PathTear	0	0	0	0
Rmt PathTear	0	0	0	0
Rmt Backup	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	
Message out-of-order	0		0	
Unknown msg-id ack	0		0	
Unknown msg-id nack	0		0	
Rcv msg-id nack	0		0	
Rcv pkt disabled interface	0		0	
Transmit failure	3		0	
Memory allocation fail	0		0	
ID allocation failed	0		0	
No path information	0		0	
Resv style conflict	0		0	
Port conflict	0		0	
Resv no interface	0		0	
PathErr to client	4		0	
ResvErr to client	0		0	
Path timeout	0		0	
Resv timeout	0		0	
No TE-link to rcv Hop	0		0	
Transmit buffer full	0		0	
P2MP RESV discarded by app	0		0	
Rate limit	0		0	
Err msg loop detected	0		0	
MP Path LP-avail rcvd	0		0	
MP Path NP-avail rcvd	0		0	
PLR bk RSB life ext	0		0	
RSB life ext for nh FRR	0		0	
MP pri PSB life ext LP	0		0	
Rcvd state rejected	0		0	
No matching senders	0		0	
Del from client	5		0	
Enhanced FRR Stats	Total		Last 5 seconds	
LP-MP LSPs retained	0		0	
NP-MP LSPs retained	0		0	

Non-MP LSPs deleted	0	0
LSPs deleted on Phop down	0	0
LSPs deleted on PPhop down	0	0
LP avail signaled LSPs	0	0
NP avail signaled LSPs	0	0
NP flag reset for Phop	0	0
LSPs retained on Cnd tear	0	0
Upstr long refresh LSPs	0	0
Upstr short refresh LSPs	0	0
Dnstr long refresh LSPs	6	0
Dnstr short refresh LSPs	0	0
PathTear ignored on MP	0	0
RRO change Remote PathTear	0	0
Primary down Rmt PathTear	0	0

show rsvp version

List of Syntax

[Syntax on page 3419](#)

[Syntax \(EX Series Switches\) on page 3419](#)

Syntax

```
show rsvp version
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show rsvp version
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.

Options

none—Display RSVP protocol settings.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

[show rsvp version on page 3422](#)

Output Fields

[Table 128 on page 3420](#) describes the output fields for the **show rsvp version** command. Output fields are listed in the approximate order in which they appear.

Table 128: show rsvp version Output Fields

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: Enabled or Disabled .
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: Aggressive , Disabled , or Normal . The default is Normal .
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.
Graceful deleting timeout	Currently configured value for the graceful-deletion-timeout statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.
NSR Mode	Status of the nonstop active routing feature for RSVP on the restarting device: Disabled , Enabled/Master , or Enabled/Standby .
NSR State	<p>State of the nonstop active routing feature for RSVP on the restarting device.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • Idle • TE-link sync complete • Neighbor sync complete • Path state sync complete • Resv state sync complete • Bypass sync complete • Init sync complete
Setup protection	Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: Enabled or Disabled .

Table 128: show rsvp version Output Fields (continued)

Field Name	Field Description
Route Session-Id count	<p>Total count of session IDs associated with the combination of all the RSVP ingress routes.</p> <p>NOTE: Starting in Junos OS Release 16.1, the show rsvp version command output displays the Route Session-Id count output field by default, irrespective of the presence of associated session IDs.</p> <p>When there are no session IDs associated with any RSVP ingress route, the <i>Route Session-Id count</i> value is zero (0).</p>
Graceful restart	Status of the graceful restart feature for RSVP on the restarting routing device: Enabled or Disabled .
Restart helper mode	Status of the helper mode feature: Enabled or Disabled . When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
Maximum helper restart time	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.
Maximum helper recovery time	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
Restart time	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
Recovery time	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
P2p transit LSP nexthop mode	Point-to-point transit LSP next-hop mode on PTX Series devices. The possible values are Chained or Unchained .
P2mp transit LSP nexthop mode	Point-to-multipoint transit LSP next-hop mode on PTX Series devices. The possible values are Chained or Unchained .

Sample Output

show rsvp version

Starting with Junos OS Release 16.1, this command also shows whether enhanced FRR procurers are enabled on the router.

```
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
  Route Session-Id count:        1
  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
  Enhanced FRR local protection: Enabled
```


traceroute mpls rsvp

Syntax

```
traceroute mpls <rsvp> lsp-name  
<detail>  
<egress>  
<exp>  
<logical-system>  
<multipoint>  
<no-resolve>  
<retries>  
<source source-address>  
<ttl>
```

Release Information

Command introduced in Junos OS Release 9.2.

egress, **multipoint**, and **ttl** options added in Junos OS Release 11.2.

Description

Trace route to a remote host for an MPLS LSP signaled by RSVP. Use **traceroute mpls rsvp** as a debugging tool to locate MPLS label-switched path (LSP) forwarding issues in a network. (Currently supported for IPv4 packets only.)

Options

lsp-name—Specify the name of the LSP to be traced.

detail—(Optional) Display detailed output.

egress—(Optional) Request that a specific point-to-multipoint egress node reply to the trace route. The trace route would follow the associated sub-LSP to the egress node.

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.

logical-system—(Optional) Specify the name of the logical system for the traceroute attempt.

multipoint—(Optional) Perform a trace route on a point-to-multipoint LSP.

no-resolve—(Optional) Specify not to resolve the hostname that corresponds to the IP address.

retries—(Optional) Specify the number of times to resend probe. The range of values is 1 through 9. The default value is 3.

source *source-address*—(Optional) Specify the source address of the outgoing traceroute packets.

ttl—(Optional) Specify the number of hops to follow before forcing the trace route to quit.

Required Privilege Level

network

List of Sample Output

[traceroute mpls rsvp on page 3425](#)

[traceroute mpls rsvp detail on page 3425](#)

[traceroute mpls rsvp multipoint \(branch node for sub-LSPs\) on page 3426](#)

[traceroute mpls rsvp multipoint \(single-hop sub-LSPs\) on page 3427](#)

Output Fields

[Table 129 on page 3424](#) 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 129: traceroute mpls rsvp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls rsvp <i>lsp-name</i> command.	all levels
ttl	Time-to-live value of the labeled packet.	none specified
Label	MPLS label used to forward the packets along the LSP.	none specified
Protocol	Signaling protocol used. For this command, it is RSVP-TE.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths). Displays Success if the trace to a hop is successful or Egress if the trace has reached the last router on the path.	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null.	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail

Table 129: traceroute mpls rsvp Output Fields (continued)

Field Name	Field Description	Level of Output
Sender timestamp	Displays the timestamp when the MPLS echo request is sent to the next hop.	detail
Receiver timestamp	Timestamp when the echo request from the previous hop is received and acknowledged with an echo response by the next hop.	detail
Response time	Time for the echo request to reach the receiver.	detail
MTU	Size of the largest packet that includes the label stack forwarded to the next hop.	detail
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none.	detail
Label stack	Label stack used to forward the packet.	detail
Path	Displays the sub-lsp path number for this traceroute, the interface used, and the destination address.	all levels

Sample Output

traceroute mpls rsvp

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta
```

```
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	299792	RSVP-TE	192.168.1.2	(null)	Success
2	299803	RSVP-TE	192.168.2.3	192.168.1.2	Success
3	3	RSVP-TE	192.168.3.4	192.168.2.3	Egress

```
Path 1 via ge-0/0/0.1 destination 127.0.0.64
```

traceroute mpls rsvp detail

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta detail
```



```
Probe options: retries 3, exp 7
```

```
Hop 192.168.1.2 Depth 1
```

```
Probe status: Success
```

```
Parent: (null)
```

```
Return code: Label-switched at stack-depth 1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 400.88 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 427.87 msec
```

```
Response time: 26.99 msec
```

```
MTU: Unknown
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299792 Protocol RSVP-TE
```

```
Hop 192.168.2.3 Depth 2
```

```
Probe status: Success
```

```
Parent: 192.168.1.2
```

```
Return code: Upstream interface index unknown label-switched at stack-depth 1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 522.13 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 548.69 msec
```

```
Response time: 26.55 msec
```

```
MTU: 1518
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299803 Protocol RSVP-TE
```

traceroute mpls rsvp multipoint (branch node for sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP where the penultimate node is a branch node for the sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
```

```
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	300000	RSVP-TE	81.1.2.2	(null)	Success
2	299968	RSVP-TE	81.2.3.3	81.1.2.2	Success
3	299952	RSVP-TE	81.3.4.4	81.2.3.3	Success
4	299920	RSVP-TE	81.4.6.6	81.3.4.4	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
4	299920	RSVP-TE	81.4.5.5	81.3.4.4	Egress

Path 2 via lt-1/2/0.102 destination 127.0.0.64

traceroute mpls rsvp multipoint (single-hop sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP with multiple single-hop sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
```

Probe options: retries 3, exp 7

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.2.2	(null)	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.8.8	(null)	Egress

Path 2 via lt-1/2/0.108 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.9.9	(null)	Egress

Path 3 via lt-1/2/0.109 destination 127.0.0.64

LDP Operational Commands

IN THIS CHAPTER

- clear ldp neighbor | 3430
- clear ldp session | 3431
- clear ldp statistics | 3433
- ping mpls ldp | 3434
- ping mpls segment routing ospf | 3437
- ping mpls segment routing isis | 3440
- show ldp database | 3443
- show ldp fec-filters | 3454
- show ldp interface | 3456
- show ldp neighbor | 3458
- show ldp overview | 3461
- show ldp p2mp tunnel | 3466
- show ldp path | 3467
- show ldp route | 3469
- show ldp session | 3480
- show ldp statistics | 3488
- show ldp traffic-statistics | 3494
- show security keychain | 3499
- traceroute mpls ldp | 3502
- traceroute mpls segment-routing ospf | 3507
- traceroute mpls segment-routing isis | 3511

clear ldp neighbor

Syntax

```
clear ldp neighbor  
<instance instance-name>  
<logical-system (all | logical-system-name)>  
<neighbor>
```

Description

Tear down Label Distribution Protocol (LDP) neighbor connections.

Options

none—Tear down connections with all LDP neighbors for all routing instances.

instance *instance name*—(Optional) Clear the LDP session for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear an LDP session for the specified neighbor (IP address) only.

Required Privilege Level

clear

RELATED DOCUMENTATION

| [show ldp neighbor](#) | 3458

List of Sample Output

[clear ldp neighbor on page 3430](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear ldp neighbor
```

```
user@host> clear ldp neighbor
```


clear ldp session

Syntax

```
clear ldp session  
<all>  
<destination>  
<instance instance-name>  
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Clear Label Distribution Protocol (LDP) sessions.

Options

all—Clear LDP sessions for all destinations for all routing instances.

destination—(Optional) Clear an LDP session for the specified destination (IP address).

instance *instance-name*—(Optional) Clear the LDP session for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

RELATED DOCUMENTATION

| [show ldp session](#) | 3480

List of Sample Output

[clear ldp session on page 3432](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ldap session

```
user@host> clear ldap session all
```


clear ldp statistics

Syntax

```
clear ldp statistics  
<instance instance-name>  
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Set all Label Distribution Protocol (LDP) statistics to zero.

Options

none—Set all LDP statistics to zero for all routing instances.

instance *instance-name*—(Optional) Clear the LDP session for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

RELATED DOCUMENTATION

[show ldp statistics | 3488](#)

[show ldp traffic-statistics | 3494](#)

List of Sample Output

[clear ldp statistics on page 3433](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear ldp statistics
```

```
user@host> clear ldp statistics
```


ping mpls ldp

Syntax

```
ping mpls ldp fec
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<p2mp root-addr ip-address lsp-id identifier>
<size bytes>
<source source-address>
<stitched-protocol>
<sweep>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

size and **sweep** options introduced in Junos OS Release 9.6.

instance option introduced in Junos OS Release 10.0.

p2mp, **root-address**, and **lsp-id** options introduced in Junos OS Release 11.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Check the operability of MPLS LDP-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.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping an LDP-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.

p2mp root-addr ip-address lsp-id identifier—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

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**).

stitched-protocol—(Optional) Protocol stitched on intermediate node.

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 User Guide*.

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.

NOTE: In a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the source address of the reply message is an interface address other than the loopback address or router ID. Starting in Junos OS Release 13.3X8, 14.2R6, 15.1R4, 15.1F6, 15.1F5-S8, 16.1R1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Required Privilege Level

network

List of Sample Output

[ping mpls ldp fec count on page 3436](#)

[ping mpls ldp p2mp root-addr lsp-id on page 3436](#)

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
```


ping mpls segment routing ospf

Syntax

```
ping mpls segment routing ospf fec
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<p2mp root-addr ip-address lsp-id identifier>
<size bytes>
<source source-address>
<stitched-protocol>
<sweep>
```

Release Information

Command introduced in Junos OS Release 19.1R1

Description

Check the operability of MPLS segment routing label-switched path (LSP) connections added by ospf protocol. Type Ctrl+c to interrupt a **ping mpls segment routing ospf** 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.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping an LDP-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.

p2mp root-addr *ip-address* *lsp-id identifier*—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

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).

stitched-protocol—(Optional) Protocol stitched on intermediate node.

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Required Privilege Level

network

List of Sample Output

[ping mpls segment-routing ospf on page 3438](#)

[ping mpls segment-routing isis on page 3438](#)

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 segment-routing ospf

```
user@host>ping mpls segment-routing ospf 6.6.6.6
```

```
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing isis

```
user@host>ping mpls segment-routing ospf 6.6.6.6 detail
```

```
Request for seq 1, to interface 333, label 22, packet size 80
Reply for seq 1, return code: Egress-ok, time: -7542.667 ms
Local transmit time: 2019-03-06 11:55:03 IST 639.914 ms
```



```
Remote receive time: 2019-03-06 11:54:56 IST 97.247 ms
Request for seq 2, to interface 333, label 22, packet size 80
Reply for seq 2, return code: Egress-ok, time: -7543.680 ms
Local transmit time: 2019-03-06 11:55:04 IST 641.965 ms
Remote receive time: 2019-03-06 11:54:57 IST 98.285 ms
Request for seq 3, to interface 333, label 22, packet size 80
Reply for seq 3, return code: Egress-ok, time: -7530.457 ms
Local transmit time: 2019-03-06 11:55:05 IST 639.923 ms
Remote receive time: 2019-03-06 11:54:58 IST 109.466 ms
Request for seq 4, to interface 333, label 22, packet size 80
Reply for seq 4, return code: Egress-ok, time: -7540.548 ms
Local transmit time: 2019-03-06 11:55:06 IST 642.870 ms
Remote receive time: 2019-03-06 11:54:59 IST 102.322 ms
Request for seq 5, to interface 333, label 22, packet size 80
Reply for seq 5, return code: Egress-ok, time: -7540.672 ms
Local transmit time: 2019-03-06 11:55:07 IST 646.870 ms
Remote receive time: 2019-03-06 11:55:00 IST 106.198 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


ping mpls segment routing isis

Syntax

```
ping mpls segment routing isis fec
<count count>
<destination address>
<detail>
<exp forwarding-class>
<instance routing-instance-name>
<logical-system (all | logical-system-name)>
<p2mp root-addr ip-address lsp-id identifier>
<size bytes>
<source source-address>
<stitched-protocol>
<sweep>
```

Release Information

Command introduced in Junos OS Release 19.1R1

Description

Check the operability of MPLS segment routing label-switched path (LSP) connections added by ISIS protocol. Type Ctrl+c to interrupt a **ping mpls segment routing isis** 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.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping an LDP-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.

p2mp root-addr *ip-address* *lsp-id identifier*—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

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).

stitched-protocol—(Optional) Protocol stitched on intermediate node.

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Required Privilege Level

network

List of Sample Output

[ping mpls segment-routing isis on page 3441](#)

[ping mpls segment-routing isis on page 3441](#)

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 segment-routing isis

```
user@host>ping mpls segment-routing isis 6.6.6.6
```

```
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls segment-routing isis

```
user@host>ping mpls segment-routing isis 6.6.6.6 detail
```

```
Request for seq 1, to interface 333, label 402006, packet size 80
Reply for seq 1, return code: Egress-ok, time: -7539.106 ms
    Local transmit time: 2019-03-06 11:54:15 IST 966.883 ms
```



```
Remote receive time: 2019-03-06 11:54:08 IST 427.777 ms
Request for seq 2, to interface 333, label 402006, packet size 80
Reply for seq 2, return code: Egress-ok, time: -7535.925 ms
Local transmit time: 2019-03-06 11:54:16 IST 978.890 ms
Remote receive time: 2019-03-06 11:54:09 IST 442.965 ms
Request for seq 3, to interface 333, label 402006, packet size 80
Reply for seq 3, return code: Egress-ok, time: -7526.134 ms
Local transmit time: 2019-03-06 11:54:17 IST 972.870 ms
Remote receive time: 2019-03-06 11:54:10 IST 446.736 ms
Request for seq 4, to interface 333, label 402006, packet size 80
Reply for seq 4, return code: Egress-ok, time: -7539.500 ms
Local transmit time: 2019-03-06 11:54:18 IST 967.014 ms
Remote receive time: 2019-03-06 11:54:11 IST 427.514 ms
Request for seq 5, to interface 333, label 402006, packet size 80
Reply for seq 5, return code: Egress-ok, time: -7539.605 ms
Local transmit time: 2019-03-06 11:54:19 IST 969.893 ms
Remote receive time: 2019-03-06 11:54:12 IST 430.288 ms

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


show ldp database

Syntax

```
show ldp database
<brief | detail | extensive>
<inet | l2circuit>
<instance instance-name>
<logical-system (all | logical-system-name)>
<p2mp>
<session session>
<summary>
```

Release Information

Command introduced before Junos OS Release 7.4.

summary option introduced in Junos OS Release 14.2.

Description

Display entries in the LDP database.

Options

none—Display standard information about all entries in the LDP database for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

inet | l2circuit—(Optional) Display only IPv4 or Layer 2 circuit bindings.

instance *instance-name*—(Optional) Display routing instance information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

p2mp—(Optional) Display point-to-multipoint binding information.

session *session*—(Optional) Display database for the specified session only. ***session*** is the destination address of the LDP session.

summary—(Optional)—Display summary output. This option displays the number of labels received and advertised for each LDP session.

Required Privilege Level

view

List of Sample Output

[show ldp database \(master\) on page 3446](#)

[show ldp database \(standby\) on page 3447](#)

[show ldp database l2circuit detail on page 3449](#)

[show ldp database l2circuit extensive on page 3449](#)

[show ldp database p2mp \(master\) on page 3450](#)

[show ldp database p2mp \(standby\) on page 3450](#)

[show ldp database session on page 3450](#)

[show ldp database \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 3451](#)

[show ldp database \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 3452](#)

[show ldp database summary on page 3453](#)

Output Fields

[Table 130 on page 3444](#) describes the output fields for the **show ldp database** command. Output fields are listed in the approximate order in which they appear.

Table 130: show ldp database Output Fields

Field Name	Field Description	Level of Output
Input label database	Label received from the other router.	All levels
Output label database	Label advertised to the other router.	All levels
<i>session-identifier</i>	Session identifier, which includes the local and remote label space identifiers.	All levels
Labels received	Number of labels received from the other router.	All levels
Labels advertised	Number of labels advertised to the other router.	All levels.
Label	Label binding to a route prefix.	All levels

Table 130: show ldp database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Prefix	<p>Route prefix.</p> <p>It can be one of the following values:</p> <ul style="list-style-type: none"> • IP prefix. • Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured. • Layer 2 encapsulation type. <p>Layer 2 encapsulation types are displayed in the format L2CKT control word status encapsulation-type vc-number, for example, L2CKT CtlfWord FRAME RELAY VC 2</p> <ul style="list-style-type: none"> • control-word-status—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> • NoCtrlWord • CtrlWord • encapsulation-type—Encapsulation type: <ul style="list-style-type: none"> • FRAME RELAY • ATM AAL5 • ATM CELL • VLAN • ETHERNET • CISCO_HDLC • PPP • VC number—Virtual circuit number. It can have any numeric value. • (Stale)—When you display the LDP database for the neighbor of a restarting router, the bindings learned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time. 	All levels
MTU	MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.	detail
VCCV Control Channel types	<p>Virtual Circuit Connection Verification (VCCV) control channel types.</p> <ul style="list-style-type: none"> • MPLS router alert label • MPLS PW label with TTL=1 	extensive

Table 130: show ldp database Output Fields (*continued*)

Field Name	Field Description	Level of Output
VCCV Control Verification types	The only valid VCCV control verification type is LSP ping .	extensive
TDM payload size	Size of the Time Division Multiplex (TDM) payload.	All levels
TDM bitrate	Bit rate for the TDM traffic.	All levels
Requested VLAN ID	(VLANs) VLAN identifier of the Layer 2 circuit.	detail
Cell bundle size	(ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.	detail
State	State of the label binding: <ul style="list-style-type: none"> • Active—Label binding has been installed and distributed appropriately. A label binding is almost always in this state. • New—New label that has not yet been distributed. <ul style="list-style-type: none"> • MapRcv—Waiting to receive a label mapping message. • MapSend—Waiting to send a label mapping message. • RelRcv—Waiting to receive a label release message. • RelRsnd—Waiting to receive a label release message before resending label mapping message. • RelSend—Waiting to send a label release message. • ReqSend—Waiting to send a label request message. • W/dSend—Waiting to send a label withdrawal message. 	detail
Age	Time elapsed since the binding was created.	detail

Sample Output

```
show ldp database (master)
```

```
user@host> show ldp database extensive
```



```

Input label database, 10.255.107.232:0--10.255.107.236:0
  Label      Prefix
299840      10.255.107.232/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
      3      10.255.107.236/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
299776      L2CKT CtrlWord VLAN VC 100
            MTU: 1500 Requested VLAN ID: 600 Flow Label T Bit: 1 Flow Label R Bit:
1
            State: Active
            Age: 9:35
            Entropy Label Capability: No
            VCCV Control Channel types:
              PWE3 control word
              MPLS router alert label
              MPLS PW label with TTL=1
            VCCV Control Verification types:
              LSP ping
              BFD with PW-ACH-encapsulation for Fault Detection
              BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 10.255.107.232:0--10.255.107.236:0
  Label      Prefix
      3      10.255.107.232/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No
299776      10.255.107.236/32
            State: Active
            Age: 9:35
            Entropy Label Capability: No

```

show ldp database (standby)

user@host> show ldp database extensive

```

Input label database, 10.255.107.236:0--10.255.107.234:0
  Label      Prefix
299808      10.255.107.230/32

```



```

State: Active
Age: 1d 2:46:36
Standby binding state:
  Map messages: 1
  Release messages: 0

```

```

Label Prefix
301136 10.255.107.232/32
State: Active
Age: 1d 2:46:36
Standby binding state:
  Map messages: 1
  Release messages: 0

```

```

Label Prefix
3 10.255.107.234/32
State: Active
Age: 1d 2:46:36
Standby binding state:
  Map messages: 1
  Release messages: 0

```

```

Label Prefix
302480 10.255.107.236/32
State: Active
Age: 1d 2:46:36
Standby binding state:
  Map messages: 1
  Release messages: 0

```

Output label database, 10.255.107.236:0--10.255.107.234:0

```

Label Prefix
299904 10.255.107.230/32
State: Active
Age: 1d 2:46:36
299936 10.255.107.232/32
State: Active
Age: 1d 2:46:36
299872 10.255.107.234/32
State: Active
Age: 1d 2:46:36
3 10.255.107.236/32
State: Active
Age: 1d 2:46:36
299952 P2MP root-addr 10.255.107.230, lsp-id 16777217
State: Active
Age: 1d 2:46:36

```


show ldp database l2circuit detail

```
user@host> show ldp database l2circuit detail
```

```
Input label database, 10.255.245.44:0--10.255.245.45:0
  Label      Prefix
  100176     L2CKT CtrlWord ATM CELL (VC Mode) VC 100
              Cell bundle size: 80
              State: Active
              Age: 9:48
  100256     L2CKT CtrlWord FRAME RELAY VC 101
              MTU: 4470
              State: Active
              Age: 9:48

Output label database, 10.255.245.44:0--10.255.245.45:0
  Label      Prefix
  100048     L2CKT CtrlWord ATM CELL (VC Mode) VC 100
              Cell bundle size: 80
              State: Active
              Age: 9:48
  100112     L2CKT CtrlWord FRAME RELAY VC 101
              MTU: 4470
              State: Active
              Age: 9:48
```

show ldp database l2circuit extensive

```
user@host> show ldp database l2circuit extensive
```

```
Input label database, 10.255.245.198:0--10.255.245.194:0
  Label      Prefix
  299872     L2CKT CtrlWord PPP VC 100
              MTU: 4470
              VCCV Control Channel types:
                  MPLS router alert label
                  MPLS PW label with TTL=1
              VCCV Control Verification types:
                  LSP ping
  Label      Prefix
              State: Active
              Age: 19:23:08
```


show ldp database p2mp (master)

```
user@host> show ldp database p2mp extensive
```

```
Input label database, 10.255.107.232:0--10.255.107.236:0
  Label      Prefix
  569649     P2MP root-addr 10.255.107.232, lsp-id 16777217
              State: Active
              Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
  Label      Prefix
  299888     P2MP root-addr 10.255.107.230, lsp-id 16777217
              State: Active
              Age: 2d 6:41:35
```

show ldp database p2mp (standby)

```
user@host> show ldp database p2mp extensive
```

```
Input label database, 10.255.107.236:0--10.255.107.232:0
  Label      Prefix
  299968     P2MP root-addr 10.255.107.230, lsp-id 16777217
              State: Active
              Age: 4d 22:21:57
  Standby binding state:
    Map messages: 1
    Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
  Label      Prefix
    3        P2MP root-addr 10.255.107.232, lsp-id 1
              State: Active
              Age: 4d 22:21:57
```

show ldp database 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
299792	1.1.1.3/32
3	1.1.1.6/32
299776	10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.6:0

Label	Prefix
299856	1.1.1.2/32
299776	1.1.1.3/32
299792	1.1.1.6/32
3	10.255.2.227/32
299888	P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7


```
299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7
```

show ldp database summary

```
user@host> show ldp database summary
```

Session ID	Labels received	Labels advertised
10.255.0.1:0--10.255.0.2:0	4	4
10.255.0.1:0--10.255.0.3:0	4	4

show ldp fec-filters

Syntax

```
show ldp fec-filters
<fec>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters.

Options

- fec**—(Optional) Display FEC filter information for the specified FEC.
- instance instance-name**—(Optional) Display FEC filter information for the specified instance.
- logical-system (all | logical-system-name)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

[show ldp fec-filters on page 3455](#)

Output Fields

[Table 131 on page 3454](#) lists the output fields for the **show ldp fec-filters** command. Output fields are listed in the approximate order in which they appear.

Table 131: show ldp fec-filters Output Fields

Field Name	Field Description
Ingress	Names of the FEC filters on the ingress routers.
Transit	Names of the FEC filters on the transit routers.

Sample Output

show ldp fec-filters

user@host> **show ldp fec-filters 10/8**

```
10.22.1.2/32
  Ingress: f1-10.22.1.2/32 (index: 3)
  Transit: (null) (index: 0)
```


show ldp interface

Syntax

```
show ldp interface
<brief | detail | extensive>
<interface-name>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Display the status of Label Distribution Protocol (LDP)-enabled interfaces.

Options

- none**—Display standard status information about all LDP-enabled interface for all routing instances.
- interface-name**—(Optional) Display information for the specified interface.
- brief | detail | extensive**—(Optional) Display the specified level of output.
- instance instance-name**—(Optional) Display information for the specified routing instance.
- logical-system (all | logical-system-name)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

[show ldp interface extensive on page 3457](#)

Output Fields

[Table 132 on page 3456](#) describes the output fields for the **show ldp interface** command. Output fields are listed in the approximate order in which they appear.

Table 132: show ldp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface name.	All levels

Table 132: show ldp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Label space ID	Label space identifier that the router is advertising on the interface.	All levels
Nbr count	Number of neighbors on the interface.	All levels
Next hello	How long until the next hello packet is sent on this interface, in seconds.	All levels
Hello interval	One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used.	detail extensive
Hold time	Configured hold time, in seconds.	detail extensive
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	extensive
Local hello interval	Locally configured hello interval.	extensive

Sample Output

show ldp interface extensive

user@host> **show ldp interface extensive**

```
Interface          Label space ID      Nbr count  Next hello
fe-0/0/3.0         10.255.245.6:0      2          0
  Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
  Local hello interval: 2, Index: 69
```


show ldp neighbor

Syntax

```
show ldp neighbor  
<brief | detail | extensive>  
<auto-targeted>  
<instance instance-name>  
<logical-system (all | logical-system-name)>  
<neighbor-address>
```

Release Information

Command introduced before Junos OS Release 7.4.

neighbor-address option added in Junos OS Release 8.5.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

auto-targeted option added in Junos OS Release 14.2.

Description

Display Label Distribution Protocol (LDP) neighbor information.

Options

none—Display standard information about LDP neighbors for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

auto-targeted—(Optional) Display information about LDP neighbors that are automatically targeted using the loopback addresses.

instance *instance-name*—(Optional) Display information for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor-address—(Optional) Display information about the specified LDP neighbor.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear ldp neighbor](#) | [3430](#)

List of Sample Output

[show ldp neighbor extensive on page 3459](#)

[show ldp neighbor auto-targeted extensive on page 3460](#)

Output Fields

[Table 133 on page 3459](#) describes the output fields for the **show ldp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 133: show ldp neighbor Output Fields

Field Name	Field Description	Level of Output
Address	IP address of the neighbor.	All levels
Interface	Interface over which the neighbor was discovered.	All levels
Label space ID	Label space identifier advertised by the neighbor.	All levels
Hold time	Remaining hold time before the neighbor expires, in seconds.	All levels
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	detail
Configuration sequence	Counter that increments whenever the neighbor changes its configuration.	detail
Up for	Length of time the LDP neighbor has been in operation.	detail extensive
Reference count	Reference count for the LDP neighbor.	extensive
Hold time	Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers.	extensive
Proposed local/peer	Hold time value proposed by the local router and the peer router.	extensive

Sample Output

show ldp neighbor extensive

```
user@host> show ldp neighbor extensive
```

```

Address          Interface      Label space ID      Hold Time
192.168.37.23    so-1/0/0.0    10.255.245.5:0      44
  Transport address: 10.255.245.5, Configuration sequence: 6
```



```
Up for 00:03:37
Reference count: 1
Hold time: 45, Proposed local/peer: 15/45
```

show ldp neighbor auto-targeted extensive

user@host> **show ldp neighbor auto-targeted extensive**

```
Address          Interface      Label space ID      Hold time
10.255.107.236    lo0.0          10.255.107.236:0    41
  Transport address: 10.255.107.236, Configuration sequence: 14
  Up for 00:10:53
  Reference count: 2
  Hold time: 45, Proposed local/peer: 45/45
  Hello interval: 15
  Hello flags: targeted
  Neighbor types: Auto-targeted
```


show ldp overview

Syntax

```
show ldp overview
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switch and QFX Series)

```
show ldp overview
<instance instance-name>
```

Release Information

Command introduced in Junos OS Release 11.2.

Description

Display LDP overview information.

Options

none— Display standard overview information about LDP for all routing instances.

instance *instance-name*— (Optional) Display LDP overview information for the specified routing instance.

logical-system (all | *logical-system-name*)— (Optional) Display LDP information from systems or a particular logical system on the devices.

Required Privilege Level

view

List of Sample Output

[show ldp overview on page 3464](#)

Output Fields

[Table 134 on page 3461](#) lists the output fields for the **show ldp overview** command. Output fields are listed in the approximate order in which they appear.

Table 134: show ldp overview Output Fields

Field Name	Field description	Level of Output
Instance	LDP routing instance.	All Levels
Router ID	Router ID of the routing device.	All Levels

Table 134: show ldp overview Output Fields (*continued*)

Field Name	Field description	Level of Output
Message ID	Unique identifier of message.	All Levels
Configuration sequence	Value of configuration sequence.	All Levels
Deaggregate	Status of control forwarding equivalence class (FEC) deaggregation on the router. By default it is disabled on the router.	All Levels
Explicit null	Advertise label 0 to the egress routing device of an LSP. Explicit null: enabled or disabled . NOTE: If you do not include the explicit-null statement in the configuration, label 3 (implicit null) is advertised.	All Levels
IPv6 tunneling	Internet Protocol version 6 tunneling: enabled or disabled .	All Levels
Strict targeted hellos	Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. Strict targeted hellos: enabled or disabled . NOTE: LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.	All Levels
Loopback if added	Loopback interface is added: yes or no .	All Levels
Route preference	Default preference value (also known as an administrative distance) assigned to each route that the routing table receives. LDP preference is: 9	All Levels
Unicast transit LSP chaining	Unicast transit LSP chaining: enabled or disabled .	All Levels
P2MP transit LSP chaining	P2MP transit LSP chaining: enabled or disabled .	All Levels
Transit LSP statistics based on route statistics	Transit LSP statistics based on route statistics: enabled or disabled .	All Levels
Longest match	Longest match for label mapping: enabled or disabled .	All Levels
Capabilities enabled	Enabled capabilities: none	All Levels

Table 134: show ldp overview Output Fields (*continued*)

Field Name	Field description	Level of Output
Timers	<ul style="list-style-type: none"> • Keepalive interval: Keepalive interval value. • Keepalive timeout: Time interval for which the neighbor LDP node waits before determining session failure. • Link hello interval: Specify how often the router sends Link Management Protocol (LMP) hello packets. • Link hello hold time: Time interval for which an LDP node waits for a hello message before declaring a neighbor is down. • Targeted hello interval: Specify how often LDP sends targeted hello messages. • Targeted hello hold time: Time interval for which a sending LSR maintains a record of targeted hello messages from the receiving LSR without receipt of another targeted hello message from that LSR. • Label withdraw delay: Time interval for withdrawing labels to reduce router workload during IGP convergence. 	All Levels
Graceful restart	<p>Graceful restart attributes.</p> <ul style="list-style-type: none"> • Restart– Graceful restart capability: enabled or disabled. • Helper– Standard graceful restart helper capability: enabled or disabled. • Restart in process– Graceful restart in process. • Reconnect time– Period of time that a restarting LSR (label switched router) designates to LDP neighbors to wait until the former reestablishes the session after restarting. • Max neighbor reconnect time– Maximum reconnect time. • Recovery time– Period of time that an LSR preserves its state across the restart. • Max neighbor recovery time– Maximum recovery time designated to LDP neighbors by a restarting LSR. 	All Levels
Traffic Engineering	<ul style="list-style-type: none"> • Bgp igp– BGP and IGP destinations: enabled or disabled. When enabled, IGP's use MPLS paths for forwarding traffic. • Both ribs– BGP and IGP destinations with routes in both RIBs: enabled or disabled. • Mpls forwarding– MPLS routes used for forwarding: enabled or disabled. 	All Levels

Table 134: show ldp overview Output Fields (*continued*)

Field Name	Field description	Level of Output
IGP	<ul style="list-style-type: none"> • Tracking igp metric—Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1). • Sync session up delay— Time interval to synchronize LDP session. 	All Levels
Session protection	<ul style="list-style-type: none"> • Session protection— Remote neighbor added to LDP configuration which enables protection for all sessions in the corresponding LDP instance: enabled or disabled. • Session protection timeout— Period of time until which the remote neighbor is connected to LSR in the absence of link neighbors. 	All Levels
Interface addresses advertising	Advertises interface address.	All Levels
Label allocation	<p>Label accounting information.</p> <ul style="list-style-type: none"> • Current number of LDP labels allocated—Number of labels currently in use. • Total number of LDP labels allocated—Cumulative number of labels being allocated. • Total number of LDP labels freed—Cumulative number of labels being freed. • Total number of LDP label allocation failures—Cumulative number of failures for allocating a label. • Current number of labels allocated by all protocols—Number of labels currently being used by routing protocols. 	All Levels

Sample Output

show ldp overview

user@host> show ldp overview

```
Instance: master
Router ID: 192.168.2.1
Message id: 0
Configuration sequence: 1
Deaggregate: disabled
```



```
Explicit null: disabled
IPv6 tunneling: disabled
Strict targeted hellos: disabled
Loopback if added: yes
Route preference: 9
Unicast transit LSP chaining: disabled
P2MP transit LSP chaining: disabled
Transit LSP statistics based on route statistics: disabled
Longest Match: enabled
Capabilities enabled: none
Timers:
  Keepalive interval: 10, Keepalive timeout: 30
  Link hello interval: 5, Link hello hold time: 15
  Targeted hello interval: 15, Targeted hello hold time: 45
  Label withdraw delay: 60
Graceful restart:
  Restart: enabled, Helper: enabled, Restart in process: false
  Reconnect time: 60000, Max neighbor reconnect time: 120000
  Recovery time: 160000, Max neighbor recovery time: 240000
Traffic Engineering:
  Bgp igp: disabled
  Both ribs: disabled
  Mpls forwarding: disabled
IGP:
  Tracking igp metric: disabled
  Sync session up delay: 10
Session protection:
  Session protection: disabled
  Session protecton timeout: 0
Interface addresses advertising:
  192.168.2.1
Label allocation:
  Current number of LDP labels allocated: 3
  Total number of LDP labels allocated: 3
  Total number of LDP labels freed: 0
  Total number of LDP label allocation failure: 0
  Current number of labels allocated by all protocols: 3
```


show ldp p2mp tunnel

Syntax

```
show ldp p2mp tunnel
<brief | detail | extensive>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced in Junos OS Release 13.3.

Description

Display LDP point-to-multipoint tunnel table information.

Options

brief | detail | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display routing instance information for the specified instance only.

logical-system (all | *logical-system-name*)—(Optional) Display LDP point-to-multipoint tunnel table information of all logical systems or a particular logical system.

Required Privilege Level

View

Sample Output

show ldp p2mp tunnel

user@host> **show ldp p2mp tunnel extensive**

```
Instance      Tunnel type      Tunnel name
0             Name             10.254.1.1:1:ldp-p2mp:mvpn:vpn-1
P2MP root-addr 10.255.107.232, lsp-id 16777217
Self id 805306372
Reference count 2
```


show ldp path

Syntax

```
show ldp path
<brief | detail | extensive>
<destination>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Display Label Distribution Protocol (LDP) label-switched paths (LSPs).

Options

- none**—Display standard information about all LDP LSPs for all routing instances.
- brief | detail | extensive**—(Optional) Display the specified level of output.
- destination**—(Optional) Restrict the output to entries that match the specified destination prefix.
- instance instance-name**—(Optional) Display information for the specified routing instance only.
- logical-system (all | logical-system-name)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

[show ldp path extensive on page 3468](#)

Output Fields

[Table 135 on page 3467](#) describes the output fields for the **show ldp path** command. Output fields are listed in the approximate order in which they appear.

Table 135: show ldp path Output Fields

Field Name	Field Description
Output Session (label)	Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.

Table 135: show ldp path Output Fields (*continued*)

Field Name	Field Description
Input Session (label)	Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.
route	MPLS route.
Attached route	Route corresponding to the LSP.
Ingress route	The router acts as the ingress for the LSP.
Reference count	Reference count for the LDP neighbor.
Transit route	Names of the forwarding equivalence class (FEC) filters on the transit routers.
Global label	MPLS label that is used globally.

Sample Output

show ldp path extensive

user@host> **show ldp path extensive**

```

Output Session (label)      Input Session (label)
10.255.14.220:0(3)          ( )
    Attached route: 10.255.14.221/32
    Reference count: 3, Global label: 3
10.255.14.220:0(100000)     10.255.14.220:0(3)
    Attached route: 10.255.14.220/32, Ingress route
    Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001)     10.255.14.220:0(100001)
    Attached route: 10.255.14.214/32, Ingress route
    Reference count: 2, Transit route, Global label: 100001

```


show ldp route

Syntax

```
show ldp route
<brief | detail | extensive>
<destination>
<fec-and-route>
<fec-only>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC).

Options

none—Display standard information about all entries in the LDP internal topology table for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

destination—(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.

fec-and-route—Display the show routes and the FECs.

fec-only—Display only LDP FECs.

instance *instance-name*—(Optional) Display entries for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

[show ldp route detail on page 3470](#)

[show ldp route extensive on page 3471](#)

[show ldp route fec-and-route on page 3472](#)

[show ldp route fec-and-route on page 3474](#)

[show ldp route fec-only on page 3477](#)

[show ldp route fec-only detail on page 3478](#)

Output Fields

Table 136 on page 3470 describes the output fields for the **show ldp route** command. Output fields are listed in the approximate order in which they appear.

Table 136: show ldp route Output Fields

Field Name	Field Description
Destination	Destination prefix.
Next-hop intf/lsp/table	Interface that is the next hop to the destination prefix.
Next-hop address	IP address of the next hop.
Session ID	LDP session ID.
Route flags	Information about the route. For example, the Ingress TTL propagate flag indicates that the time-to-live (TTL) value is being propagated with the route.
Bound to outgoing label	The route has been bound to LSPs with the label being distributed for that LSP.
Topology entry	The topology that the route is bound to.
Ingress route status	Status of the ingress route. For example, it could be Active or Inactive .
Last modified	The length of time since the ingress route status last changed.
Last event(s)	The last event that occurred.

Sample Output

show ldp route detail

```
user@host> show ldp route 10.255.8.5 detail
```

```

Destination      Next-hop intf/lsp      Next-hop address
10.255.8.5/32    f1
  Session ID 10.255.170.84:0--10.255.170.92:0
                fe-0/0/0.0          192.168.100.2
```



```

Session ID 10.255.170.84:0--10.255.8.5:0
          so-0/2/1.0
Session ID 10.255.170.84:0--10.255.8.5:0
          so-0/2/2.0
Session ID 10.255.170.84:0--10.255.8.3:0
Bound to outgoing label 299776, Topology entry: 0x8c38a80
BFD dest addr  BFD state LSP-ping Next-hop addr  Next-hop intf/lsp
127.0.0.64      up        up        192.168.100.2  fe-0/0/0.0
127.0.1.64      up        up                so-0/2/1.0
127.0.2.64      up        up                so-0/2/2.0
127.0.3.64      up        up                fl
.....

```

show ldp route extensive

user@host> show ldp route extensive

```

Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.0/30       ge-1/2/0.18             10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.4/30       ge-1/2/0.18             10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.8/30       ge-1/2/1.21             10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.12/30      ge-1/2/1.21             10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.16/30      ge-1/2/0.18
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.18/32      ge-1/2/0.18
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address
10.0.0.20/30      ge-1/2/1.21
  Route flags: None
Destination      Next-hop intf/lsp/table  Next-hop address

```



```

10.0.0.21/32
  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 route fec-and-route

user@host> show ldp route fec-and-route

```

Destination      Next-hop intf/lsp/table      Next-hop address
10.4.0.0/16       fxp0.0                      10.92.31.254

```


10.5.0.0/16	fxp0.0	10.92.31.254
10.6.128.0/17	fxp0.0	10.92.31.254
10.9.0.0/16	fxp0.0	10.92.31.254
10.10.0.0/16	fxp0.0	10.92.31.254
10.13.4.0/23	fxp0.0	10.92.31.254
10.13.10.0/23	fxp0.0	10.92.31.254
10.82.0.0/15	fxp0.0	10.92.31.254
10.84.0.0/16	fxp0.0	10.92.31.254
10.85.12.0/22	fxp0.0	10.92.31.254
10.92.0.0/16	fxp0.0	10.92.31.254
10.92.16.0/20	fxp0.0	
10.92.20.175/32		
10.94.0.0/16	fxp0.0	10.92.31.254
10.99.0.0/16	fxp0.0	10.92.31.254
10.102.0.0/16	fxp0.0	10.92.31.254
10.150.0.0/16	fxp0.0	10.92.31.254
10.155.0.0/16	fxp0.0	10.92.31.254
10.157.64.0/19	fxp0.0	10.92.31.254
10.160.0.0/16	fxp0.0	10.92.31.254
10.204.0.0/16	fxp0.0	10.92.31.254
10.205.0.0/16	fxp0.0	10.92.31.254
10.206.0.0/16	fxp0.0	10.92.31.254
10.207.0.0/16	fxp0.0	10.92.31.254
10.209.0.0/16	fxp0.0	10.92.31.254
10.212.0.0/16	fxp0.0	10.92.31.254
10.213.0.0/16	fxp0.0	10.92.31.254
10.214.0.0/16	fxp0.0	10.92.31.254
10.215.0.0/16	fxp0.0	10.92.31.254
10.216.0.0/16	fxp0.0	10.92.31.254
10.218.13.0/24	fxp0.0	10.92.31.254
10.218.14.0/24	fxp0.0	10.92.31.254
10.218.16.0/20	fxp0.0	10.92.31.254
10.218.32.0/20	fxp0.0	10.92.31.254
10.227.0.0/16	fxp0.0	10.92.31.254
10.255.111.0/24	ge-0/0/2.0	11.11.11.2
10.255.111.1/32	ge-0/0/2.0	11.11.11.2
10.255.111.2/32	ge-0/0/2.0	11.11.11.2
10.255.111.3/32	ge-0/0/2.0	11.11.11.2
10.255.111.4/32	ge-0/0/2.0	11.11.11.2
10.255.111.4/32	ge-0/0/2.0	11.11.11.2
10.255.112.1/32	lo0.0	
10.255.112.1/32	lo0.0	
10.255.112.2/32	ge-0/0/2.0	11.11.11.2
10.255.112.2/32	ge-0/0/2.0	11.11.11.2

11.11.11.0/24	ge-0/0/2.0	
11.11.11.1/32		
12.12.12.0/24	ge-0/0/2.0	11.11.11.2
15.15.15.0/24	ge-0/0/1.0	
15.15.15.1/32		
22.22.22.0/24	ge-0/0/0.0	
22.22.22.1/32		
23.23.23.0/24	ge-0/0/2.0	11.11.11.2
24.24.24.0/24	ge-0/0/2.0	11.11.11.2
25.25.25.0/24	ge-0/0/2.0	11.11.11.2
128.92.17.45/32	ge-0/0/2.0	11.11.11.2
128.92.20.175/32	lo0.0	
128.92.21.186/32	ge-0/0/2.0	11.11.11.2
128.92.25.135/32	ge-0/0/2.0	11.11.11.2
128.92.27.91/32	ge-0/0/2.0	11.11.11.2
128.92.28.70/32	ge-0/0/2.0	11.11.11.2
172.16.0.0/12	fxp0.0	10.92.31.254
192.168.0.0/16	fxp0.0	10.92.31.254
192.168.102.0/23	fxp0.0	10.92.31.254
207.17.136.0/24	fxp0.0	10.92.31.254
207.17.136.192/32	fxp0.0	10.92.31.254
207.17.137.0/24	fxp0.0	10.92.31.254
224.0.0.5/32		

show ldp route fec-and-route

user@host> show ldp route fec-and-route

Destination	Next-hop intf/lsp/table	Next-hop address
10.4.0.0/16	fxp0.0	10.92.31.254
10.5.0.0/16	fxp0.0	10.92.31.254
10.6.128.0/17	fxp0.0	10.92.31.254
10.9.0.0/16	fxp0.0	10.92.31.254
10.10.0.0/16	fxp0.0	10.92.31.254
10.13.4.0/23	fxp0.0	10.92.31.254
10.13.10.0/23	fxp0.0	10.92.31.254
10.82.0.0/15	fxp0.0	10.92.31.254
10.84.0.0/16	fxp0.0	10.92.31.254
10.85.12.0/22	fxp0.0	10.92.31.254
10.92.0.0/16	fxp0.0	10.92.31.254
10.92.16.0/20	fxp0.0	
10.92.20.175/32		
10.94.0.0/16	fxp0.0	10.92.31.254
10.99.0.0/16	fxp0.0	10.92.31.254


```

10.102.0.0/16      fxp0.0          10.92.31.254
10.150.0.0/16      fxp0.0          10.92.31.254
10.155.0.0/16      fxp0.0          10.92.31.254
10.157.64.0/19     fxp0.0          10.92.31.254
10.160.0.0/16      fxp0.0          10.92.31.254
10.204.0.0/16      fxp0.0          10.92.31.254
10.205.0.0/16      fxp0.0          10.92.31.254
10.206.0.0/16      fxp0.0          10.92.31.254
10.207.0.0/16      fxp0.0          10.92.31.254
10.209.0.0/16      fxp0.0          10.92.31.254
10.212.0.0/16      fxp0.0          10.92.31.254
10.213.0.0/16      fxp0.0          10.92.31.254
10.214.0.0/16      fxp0.0          10.92.31.254
10.215.0.0/16      fxp0.0          10.92.31.254
10.216.0.0/16      fxp0.0          10.92.31.254
10.218.13.0/24     fxp0.0          10.92.31.254
10.218.14.0/24     fxp0.0          10.92.31.254
10.218.16.0/20     fxp0.0          10.92.31.254
10.218.32.0/20     fxp0.0          10.92.31.254
10.227.0.0/16      fxp0.0          10.92.31.254
10.255.111.0/24    ge-0/0/2.0      11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
10.255.111.1/32    ge-0/0/2.0      11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300192, Topology entry: 0xb5de1b0
  Ingress route status: Active, Last modified: 09:57:49 ago
  Last event(s): Rebind
  Route flags: Transit TTL propagate, Allow longest match
Destination        Next-hop intf/lsp/table      Next-hop address
10.255.111.2/32    ge-0/0/2.0                11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300208, Topology entry: 0xb5de1f8
  Ingress route status: Active, Last modified: 09:57:49 ago
  Last event(s): Rebind
  Route flags: Transit TTL propagate, Allow longest match
Destination        Next-hop intf/lsp/table      Next-hop address
10.255.111.3/32    ge-0/0/2.0                11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300224, Topology entry: 0xb5de240
  Ingress route status: Active, Last modified: 09:57:49 ago
  Last event(s): Rebind
  Route flags: Transit TTL propagate, Allow longest match
Destination        Next-hop intf/lsp/table      Next-hop address
10.255.111.4/32    ge-0/0/2.0                11.11.11.2

```



```

Session ID 10.255.112.1:0--10.255.112.2:0
Bound to outgoing label 300112, Topology entry: 0xb5de708
Ingress route status: Active, Last modified: 10:10:56 ago
Last event(s): Evaluate Update ingress route Update transit route
Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination      Next-hop intf/lsp/table      Next-hop address
10.255.111.4/32   ge-0/0/2.0                   11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300112, Topology entry: 0xb5de708
  Ingress route status: Active, Last modified: 10:10:56 ago
  Last event(s): Evaluate Update ingress route Update transit route
  Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination      Next-hop intf/lsp/table      Next-hop address
10.255.112.1/32   lo0.0
  Bound to outgoing label 3, Topology entry: 0xb5de120
  Ingress route status: Inactive
  Last event(s): Evaluate Update history
  Route type: Egress route
  Route flags: Allow longest match
Destination      Next-hop intf/lsp/table      Next-hop address
10.255.112.1/32   lo0.0
  Bound to outgoing label 3, Topology entry: 0xb5de120
  Ingress route status: Inactive
  Last event(s): Evaluate Update history
  Route type: Egress route
  Route flags: Allow longest match
Destination      Next-hop intf/lsp/table      Next-hop address
10.255.112.2/32   ge-0/0/2.0                   11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300064, Topology entry: 0xb5de630
  Ingress route status: Active, Last modified: 10:11:04 ago
  Last event(s): Update ingress route
  Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination      Next-hop intf/lsp/table      Next-hop address
10.255.112.2/32   ge-0/0/2.0                   11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300064, Topology entry: 0xb5de630
  Ingress route status: Active, Last modified: 10:11:04 ago
  Last event(s): Update ingress route
  Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination      Next-hop intf/lsp/table      Next-hop address
11.11.11.0/24     ge-0/0/2.0
11.11.11.1/32
12.12.12.0/24     ge-0/0/2.0                   11.11.11.2

```



```

    Session ID 10.255.112.1:0--10.255.112.2:0
15.15.15.0/24      ge-0/0/1.0
15.15.15.1/32
22.22.22.0/24      ge-0/0/0.0
22.22.22.1/32
23.23.23.0/24      ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
24.24.24.0/24      ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
25.25.25.0/24      ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
128.92.17.45/32     ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
128.92.20.175/32    lo0.0
128.92.21.186/32    ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
128.92.25.135/32    ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
128.92.27.91/32     ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
128.92.28.70/32     ge-0/0/2.0      11.11.11.2
    Session ID 10.255.112.1:0--10.255.112.2:0
172.16.0.0/12       fxp0.0          10.92.31.254
192.168.0.0/16       fxp0.0          10.92.31.254
192.168.102.0/23     fxp0.0          10.92.31.254
207.17.136.0/24      fxp0.0          10.92.31.254
207.17.136.192/32    fxp0.0          10.92.31.254
207.17.137.0/24      fxp0.0          10.92.31.254
224.0.0.5/32

```

show ldp route fec-only

user@host> show ldp route fec-only

```

user@host_re0> show ldp route fec-only
Destination      Next-hop intf/lsp/table      Next-hop address
10.255.111.1/32   ge-0/0/2.0                  11.11.11.2
10.255.111.2/32   ge-0/0/2.0                  11.11.11.2
10.255.111.3/32   ge-0/0/2.0                  11.11.11.2
10.255.111.4/32   ge-0/0/2.0                  11.11.11.2
10.255.112.1/32   lo0.0
10.255.112.2/32   ge-0/0/2.0                  11.11.11.2

```


show ldp route fec-only detail

```
user@host> show ldp route fec-only detail
```

```

Destination          Next-hop intf/lsp/table      Next-hop address
10.255.111.1/32      ge-0/0/2.0                  11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300192, Topology entry: 0xb5de1b0
  Ingress route status: Active, Last modified: 09:55:10 ago
  Last event(s): Rebind
  Route flags: Transit TTL propagate, Allow longest match
Destination          Next-hop intf/lsp/table      Next-hop address
10.255.111.2/32      ge-0/0/2.0                  11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300208, Topology entry: 0xb5de1f8
  Ingress route status: Active, Last modified: 09:55:10 ago
  Last event(s): Rebind
  Route flags: Transit TTL propagate, Allow longest match
Destination          Next-hop intf/lsp/table      Next-hop address
10.255.111.3/32      ge-0/0/2.0                  11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300224, Topology entry: 0xb5de240
  Ingress route status: Active, Last modified: 09:55:10 ago
  Last event(s): Rebind
  Route flags: Transit TTL propagate, Allow longest match
Destination          Next-hop intf/lsp/table      Next-hop address
10.255.111.4/32      ge-0/0/2.0                  11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300112, Topology entry: 0xb5de708
  Ingress route status: Active, Last modified: 10:08:17 ago
  Last event(s): Evaluate Update ingress route Update transit route
  Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match
Destination          Next-hop intf/lsp/table      Next-hop address
10.255.112.1/32      lo0.0
  Bound to outgoing label 3, Topology entry: 0xb5de120
  Ingress route status: Inactive
  Last event(s): Evaluate Update history
  Route type: Egress route
  Route flags: Allow longest match
Destination          Next-hop intf/lsp/table      Next-hop address
10.255.112.2/32      ge-0/0/2.0                  11.11.11.2
  Session ID 10.255.112.1:0--10.255.112.2:0
  Bound to outgoing label 300064, Topology entry: 0xb5de630
  Ingress route status: Active, Last modified: 10:08:25 ago

```


Last event(s): Update ingress route

Route flags: Ingress TTL propagate, Transit TTL propagate, Allow longest match

show ldp session

Syntax

```
show ldp session
<brief | detail | extensive>
<auto-targeted>
<destination>
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

auto-targeted option added in Junos OS Release 14.2.

Description

Display information about Label Distribution Protocol (LDP) sessions.

Options

none—Display standard information about all LDP sessions for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

auto-targeted—(Optional) Display information about LDP sessions that are automatically targeted using loopback addresses.

destination—(Optional) Restrict LDP session display to the specified address.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, 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.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear ldp session](#) | [3431](#)

List of Sample Output

[show ldp session brief on page 3484](#)

[show ldp session detail on page 3485](#)

[show ldp session extensive on page 3485](#)

[show ldp session auto-targeted detail on page 3486](#)

Output Fields

Table 137 on page 3481 describes the output fields for the **show ldp session** command. Output fields are listed in the approximate order in which they appear.

Table 137: show ldp session Output Fields

Field Name	Field Description	Level of Output
Address	Transport address of the session.	any
State	State of the session: Nonexistent , Connecting , Initialized , OpenRec , OpenSent , Operational , or Closing . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt.	any
Connection	TCP connection state: Closed , Opening , or Open .	any
Hold time	Time remaining until the session will be closed, in seconds.	any
Session ID	LDP identifiers of the peers of this session.	detail extensive
Next keepalive	Time until next keepalive is sent, in seconds.	detail extensive
Active	Whether the local router is playing the active role in the session and during session establishment.	detail extensive
Passive	Whether the local router is playing the passive role in the session and during session establishment.	detail extensive
Maximum PDU	Maximum protocol data unit (PDU) size (packet size) for the session.	detail extensive
Hold time	Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the keepalive-timeout statement configured at the [edit protocols ldp] hierarchy level.	detail extensive
Neighbor count	Number of neighbors that are contributing to the session.	detail extensive
Neighbor types	Category of LDP session: discovered or auto-targeted .	any
Keepalive interval	Keepalive interval, in seconds.	detail extensive

Table 137: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
Reason	Reason the session went down: <ul style="list-style-type: none"> • Aborted graceful restart • Authentication key was changed • Bad type length value (TLV) • Bad protocol data unit (PDU) packets • Command-line interface (CLI) command • Connect time expired • Connection error • Connection reset • Error during initialization • Hold time expired • No adjacency or all adjacencies down • Notification received • Received notification from peer • Unexpected End of File (EOF) • Unknown reason 	detail extensive
Number of session flaps	Number of times the session changes from up to down.	detail extensive
Restarting	LDP is in the process of gracefully restarting.	detail extensive
Capabilities advertised	LDP capabilities advertised to a peer.	detail extensive
Capabilities received	LDP capabilities received from a peer.	detail extensive

Table 137: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protection	Information about the status of MPLS LDP session protection.	detail extensive
restart complete in <i>nnn</i> msec	Amount of time (in milliseconds) remaining until graceful restart is declared complete.	detail extensive
Authentication type	Shows the longest match MD5 authentication	detail extensive
Local	<p>Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> • Restart—Status of the graceful restart feature at the local end of the LDP session: enabled or disabled. • Helper mode—Status of the helper mode feature at the local end of the LDP session: enabled or disabled. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures. • Reconnect time—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is 60000 msec and is not configurable. (Reconnect timeout refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.) 	detail extensive
Remote	<p>Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> • Restart—Status of the graceful restart feature at the remote end of the LDP session: enabled or disabled. • Helper mode—Status of the helper mode feature at the remote end of the LDP session: enabled or disabled. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures. • Reconnect time—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors. 	detail extensive
Local maximum recovery time	Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).	detail extensive
Next-hop addresses received	Next-hop addresses received on the session.	detail extensive

Table 137: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Queue depth	Number of messages that are queued for sending to the peers in the group.	extensive
Message type	<p>Type of message being sent:</p> <ul style="list-style-type: none"> • Initialization—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established. • Keepalive—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them. • Notification—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer. • Address—Message sent by an LSR to an LDP peer to advertise interface addresses. • Address withdraw—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address. • Label mapping—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC). • Label request—Message sent by an LSR to an LDP peer to request a label mapping for an FEC. • Label withdraw—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping. • Label release—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released. • Label abort—Message sent by an LSR to an LDP peer to abort a label request message. • Total—Messages sent and received during the lifetime of the session. • Last 5 seconds—Messages sent and received during the current session. 	extensive

Sample Output

```
show ldp session brief
```

```
user@host> show ldp session brief
```


Address	State	Connection	Hold time
10.255.72.160	Operational	Open	21
10.255.72.164	Operational	Open	20
10.255.72.172	Operational	Open	21

show ldp session detail

user@host> show ldp session detail

```

Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
  Session ID: 192.168.0.2:0--192.168.0.3:0
  Next keepalive in 7 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.0.2, Remote address: 192.168.0.3
  Up for 00:00:02
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream unsolicited
  Negotiated Label Advertisement mode: Downstream unsolicited
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    10.0.0.5
    10.0.0.33

```

show ldp session extensive

user@host> show ldp session extensive

```

Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
  Session ID: 192.168.0.2:0--192.168.0.3:0
  Next keepalive in 2 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.0.2, Remote address: 192.168.0.3

```



```

Up for 00:05:37
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
Queue depth: 0

```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Initialization	1	1	0	0
Keepalive	33	33	1	1
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	7	5	0	0
Label request	0	0	0	0
Label withdraw	3	1	0	0
Label release	1	3	0	0
Label abort	0	0	0	0

show ldp session auto-targeted detail

```
user@host> show ldp session auto-generated detail
```

```

Address: 192.168.1.5, State: Operational, Connection: Open, Hold time: 25
  Session ID: 192.168.1.1:0--192.168.1.5:0
  Next keepalive in 5 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered, Auto-targeted
                    ^^^^^^^^^^^^^^^
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.1.1, Remote address: 192.168.1.5
  Up for 00:00:34
  Capabilities advertised: none
  Capabilities received: none

```



```
Protection: disabled
Local - Restart: disabled, Helper mode: enabled
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  192.168.1.2
  192.168.1.3
```


show ldp statistics

Syntax

```
show ldp statistics
<instance instance-name>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Display Label Distribution Protocol (LDP) statistics.

Options

- none**—Display LDP statistics for all routing instances.
- instance *instance-name***—(Optional) Display information for the specified routing instance only.
- logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

| [clear ldp statistics](#) | [3433](#)

List of Sample Output

[show ldp statistics on page 3492](#)

Output Fields

[Table 138 on page 3488](#) lists the output fields for the **show ldp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 138: show ldp statistics Output Fields

Field Name	Field Description
Total Sent, Received	Total number of each message type sent and received.

Table 138: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
Last 5 seconds Sent, Received	Number of each message type sent and received in the last 5 seconds.
Message type	<p>LDP message types:</p> <ul style="list-style-type: none"> • Hello—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor. • Initialization—Messages that indicate an LDP session has started. • Keepalive—Messages that ensure that the keepalive timeout is not exceeded. • Notification—Advisory information and signal error information. • Address—Messages with address information. • Address withdrawal—Messages regarding address withdrawal. • Label mapping—Messages with label mapping information. • Label request—Request for a label mapping from a neighboring router. • Label withdrawal—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • Label release—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • Label abort—Messages about label interruptions. • All UDP—All hello messages sent by LSRs to the well-known UDP port, 646. • All TCP—All LDP session messages.

Table 138: show ldp statistics Output Fields *(continued)*

Field Name	Field Description
Event type	

Table 138: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
	<p>LDP events and errors:</p> <ul style="list-style-type: none"> • Sessions opened—Number of LDP sessions that have been opened. • Sessions closed—Number of LDP sessions that have been closed. • Topology changes—Number of changes to the known LDP topology. • No interface—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages. • No session—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers. • No adjacency—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected. • Unknown version—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. • Malformed PDU—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session. <p>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.</p> <p>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).</p> <ul style="list-style-type: none"> • Malformed message—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them. <p>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.</p> <p>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.</p> <p>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.</p> <p>The LDP message is considered to be malformed if the message is missing one or more mandatory parameters.</p> <ul style="list-style-type: none"> • Unknown message type—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message. • Inappropriate message—The message is not of the type that the receiver expects to receive. • Malformed TLV—The TLV ILength is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR.

Table 138: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> • Bad TLV value—The TLV Length is too large. • Missing TLV—The TLV is missing one or more mandatory parameters. • PDU too large—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined.
Total	Total number of each event or error.
Last 5 seconds	Number of each event or error in the last 5 seconds.

Sample Output

show ldp statistics

user@host> show ldp statistics

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	265	263	2	2
Initialization	2	2	0	0
Keepalive	112	111	1	0
Notification	0	0	0	0
Address	2	2	0	0
Address withdraw	0	0	0	0
Label mapping	7	6	0	0
Label request	0	0	0	0
Label withdraw	2	0	0	0
Label release	0	2	0	0
Label abort	0	0	0	0
All UDP	265	263	2	2
All TCP	123	121	1	0
Event type		Total	Last 5 seconds	
Sessions opened		2	0	
Sessions closed		0	0	
Topology changes		11	0	
No interface		0	0	

No session	0	0
No adjacency	0	0
Unknown version	0	0
Malformed PDU	0	0
Malformed message	0	0
Unknown message type	0	0
Inappropriate message	0	0
Malformed TLV	0	0
Bad TLV value	0	0
Missing TLV	0	0
PDU too large	0	0

show ldp traffic-statistics

Syntax

```
show ldp traffic-statistics  
<instance instance-name>  
<logical-system (all | logical-system-name)>  
<p2mp>
```

Release Information

Command introduced before Junos OS Release 7.4.

p2mp option added in Junos OS Release 11.2.

Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description

Display Label Distribution Protocol (LDP) traffic statistics.

NOTE: If nonstop active routing features is configured, **show ldp traffic-statistics** command is not supported on backup Routing Engines.

Options

none—Display LDP traffic statistics for all routing instances.

instance *instance-name*—(Optional) Display LDP traffic statistics for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

p2mp—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.

Additional Information

To collect output from this command on a periodic basis, configure the [traffic-statistics](#) statement for the LDP protocol. For more information, see the *Junos MPLS Applications Configuration Guide*.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear ldp statistics](#) | [3433](#)

[Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain | 1024](#)

[Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs | 1082](#)

List of Sample Output

[show ldp traffic-statistics on page 3496](#)

[show ldp traffic-statistics p2mp \(Ingress or transit router only, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 3497](#)

[show ldp traffic-statistics p2mp \(Multipoint LDP with Multicast-Only Fast Reroute\) on page 3498](#)

Output Fields

Table 139 on page 3495 lists the output fields for the **show ldp traffic-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 139: show ldp traffic-statistics Output Fields

Field Name	Field Description
Message type	LDP message types.
FEC	Forwarding equivalence class (FEC) for which LDP traffic statistics are collected. For P2MP LSPs, FEC appears as a combination of root address and the LSP ID (root_addr:lsp_id). For M-LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address (root_addr:lsp_id/grp,src).
Type	Type of traffic originating from a router, either Ingress (originating from this router) or Transit (forwarded through this router).
Packets	Number of packets passed by the FEC since its LSP came up.
Bytes	Number of bytes of data passed by the FEC since its LSP came up.
Shared	Whether a label is shared by prefixes: Yes or No . A Yes value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
Nexthop	The next hop address for P2MP LSPs. (This is the downstream LDP Session ID.)

Table 139: show ldp traffic-statistics Output Fields (*continued*)

Field Name	Field Description
Label	<p>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.</p> <p>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.</p>
Backup route	For multipoint LDP with MoFRR, the route that is used if the primary route becomes unavailable.

Sample Output

show ldp traffic-statistics

user@host> **show ldp traffic-statistics**

FEC	Type	Packets	Bytes	Shared
10.35.3.0/30	Transit	0	0	Yes
	Ingress	0	0	No
10.35.10.1/32	Transit	0	0	Yes
	Ingress	0	0	No
10.255.245.214/32	Transit	0	0	No
	Ingress	11	752	No

192.168.37.36/30	Transit	0	0	Yes
	Ingress	0	0	No
FEC(root_addr:lsp_id) Nexthop Packets Bytes Shared				
10.255.72.160:16777217	192.168.8.81	152056	14597376	No
	192.168.8.1	152056	14597376	No
	192.168.8.65	152056	14597376	No
NET FEC Statistics:				
FEC	Type	Packets	Bytes	Shared
10.255.107.230/32	Transit	30858	2022345	No
	Ingress	20	5120	No

show ldp traffic-statistics p2mp (Ingress or transit router only, Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

user@host> **show ldp traffic-statistics p2mp**

P2MP FEC Statistics:				
FEC(root_addr:lsp_id/grp,src)	Nexthop	Packets	Bytes	
Shared				
11.99.0.73:239.10.0.1,11.98.0.10	11.99.0.117	243408	121217184	
No				
	11.99.0.13	236286	117670428	
No				
11.99.0.73:239.10.0.2,11.98.0.10	11.99.0.117	248800	123902400	
No				
	11.99.0.13	240759	119897982	
No				
11.99.0.73:239.10.0.1,11.98.0.20	11.99.0.117	250286	124642428	
No				
	11.99.0.13	243741	121383018	
No				
11.99.0.73:239.10.0.2,11.98.0.20	11.99.0.117	252970	125979060	
No				
	11.99.0.13	245218	122118564	
No				

show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute)

user@host> show ldp traffic-statistics p2mp

P2MP FEC Statistics:			
FEC(root_addr:lsp_id/grp,src)	Nexthop	Packets	Bytes
Shared			
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301568	1.3.8.2	0	0
No	1.3.4.2	0	0
No			
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route	1.3.4.2	0	0
No	1.3.8.2	0	0
No			
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301600	1.3.8.2	0	0
No	1.3.4.2	0	0
No			
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route	1.3.4.2	0	0
No	1.3.8.2	0	0
No			

show security keychain

Syntax

```
show security keychain  
<brief | detail>
```

Release Information

Command introduced in Junos OS Release 11.2.

Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description

Display information about authentication keychains configured for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.

Options

none—Display information about authentication keychains.

brief | detail—(Optional) Display the specified level of output.

Required Privilege Level

view

List of Sample Output

[show security keychain brief on page 3501](#)

[show security keychain detail on page 3501](#)

Output Fields

[Table 140 on page 3499](#) describes the output fields for the **show security keychain** command. Output fields are listed in the approximate order in which they appear.

Table 140: show security keychain Output Fields

Field Name	Field Description	Level of Output
keychain	The name of the keychain in operation.	All levels
Active-ID Send	Number of routing protocols packets sent with the active key.	All levels
Active-ID Receive	Number of routing protocols packets received with the active key.	All levels
Next-ID Send	Number of routing protocols packets sent with the next key.	All levels

Table 140: show security keychain Output Fields *(continued)*

Field Name	Field Description	Level of Output
Next-ID Receive	Number of routing protocols packets received with the next key.	All levels
Transition	Amount of time until the current key will be replaced with the next key in the keychain.	All levels
Tolerance	Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.	All levels
Id	Identification number configured for the current key.	detail
Algorithm	Authentication algorithm configured for the current key.	detail
State	<p>State of the current key.</p> <p>The value can be:</p> <ul style="list-style-type: none"> • receive • send • send-receive <p>For the active key, the State can be send-receive, send, or receive. For keys that have a future start time, the State is inactive. Compare the State field to the Mode field.</p>	detail
Option	<p>For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets.</p> <p>The values can be:</p> <ul style="list-style-type: none"> • basic—Based on RFC 5304. • isis-enhanced—Based on RFC 5310. <p>The default value is basic. When you configure the isis-enhanced option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.</p> <p>When you configure basic (or do not include the options statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.</p> <p>Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.</p>	detail

Table 140: show security keychain Output Fields (continued)

Field Name	Field Description	Level of Output
Start-time	Time that the current key became active.	detail
Mode	<p>Mode of each key (Informational only.)</p> <p>The value can be</p> <ul style="list-style-type: none"> ● receive ● send ● send-receive <p>The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the Mode can be send-receive, send, or receive, regardless of the configured start-time. Compare the Mode field to the State field.</p>	detail

Sample Output

show security keychain brief

```
user@host> show security keychain brief
```

keychain	Active-ID		Next-ID		Transition	Tolerance
	Send	Receive	Send	Receive		
hakr	3	3	1	1	1d 23:58	3600

show security keychain detail

```
user@host> show security keychain detail
```

keychain	Active-ID		Next-ID		Transition	Tolerance
	Send	Receive	Send	Receive		
hakr	3	3	1	1	1d 23:58	3600
Id 3, Algorithm hmac-md5, State send-receive, Option basic						
Start-time Wed Aug 11 16:28:00 2010, Mode send-receive						
Id 1, Algorithm hmac-md5, State inactive, Option basic						
Start-time Fri Aug 20 11:30:57 2010, Mode send-receive						

traceroute mpls ldp

Syntax

```
traceroute mpls <ldp> fec
<destination ip-address>
<detail>
<exp exp>
<fanout fanout-number>
<logical-system logical-system-name>
<no-resolve>
<paths maximum-paths>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source ip-address>
<ttl value>
<update>
<wait seconds>
```

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 ip-address—(Optional) Specify the destination address to use when sending probes.

Values: The destination IP address must be within the 127.0.0.0/8 IP address space for Operation, Administration, and Maintenance (OAM) packets.

detail—(Optional) Display detailed output.

exp exp—(Optional) Specify the class-of-service to use when sending probes.

Range: 0 through 7

Default: 7

fanout *fanout-number*—(Optional) Specify the maximum number of nexthops to search per node.

Range: 1 through 16

Default: 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 *maximum-paths*—(Optional) Specify the maximum number of paths to search.

Range: 1 through 255

Default: 16

pipe-mode—(Optional) Specify to trace only nodes that understand LDP FEC.

In an interoperation with other vendor devices or devices running Junos OS Release that do not support tracing of hierarchical LSPs as described in RFC 6424, continuous non-complaint probe status is displayed in the **traceroute mpls ldp** command output. To avoid this LDP loop creation, use the **pipe-mode** option with the **traceroute mpls ldp fec** command.

NOTE: Even after using the **traceroute mpls ldp fec pipe-mode** command, one or more intermediate transit nodes that do not understand LDP FEC can return non-complaint probe status in the command output.

retries *retries-number*—(Optional) Specify the number of times to resend probe values.

Range: 1 through 9

Default: 3

routing-instance *routing-instance-name*—(Optional) Specify the name of the routing instance for the 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.

Range: 1 through 125

Default: 64

update—(Optional) Update database contents with traceroute results.

wait *seconds*—(Optional) Specify the number of seconds to wait before resending a probe.

Range: 5 through 15

Default: 10

Required Privilege Level

network

List of Sample Output

[traceroute mpls ldp on page 3505](#)

[traceroute mpls ldp detail on page 3505](#)

Output Fields

Table 119 on page 3350 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 141: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls ldp fec command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null .	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null .	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail
Response time	Time for the echo request to reach the receiver.	detail
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none .	detail

Table 141: traceroute mpls ldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Label Stack	Label stack used to forward the packet.	detail

Sample Output

traceroute mpls ldp

```
user@router> traceroute mpls ldp 4.4.4.4
```

```

Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
ttl   Label  Protocol  Address      Previous Hop  Probe Status
  1   100016  LDP       24.24.24.1   (null)        Success
  2   100000  LDP       20.20.20.2   24.24.24.1    Success
  3     3     LDP       22.22.22.4   20.20.20.2    Egress

Path 1 via fe-0/3/3.101 destination 127.0.0.64
```

traceroute mpls ldp detail

```
user@router> traceroute mpls ldp 4.4.4.4 detail
```

```

Probe Options: ttl 64, retries 3, wait 10, paths 3, exp 7
Hop 24.24.24.1 Depth 1
  Parent (null)
  Return code: Label switched at stack-depth 1
  Response time 165.93 msec
  Multipath type: IP bitmask
    Address Range 1: 127.0.0.0 ~ 127.0.3.255
  Label Stack:
    Label 1 Value 100032 Protocol LDP

Hop 20.20.20.2 Depth 2
  Parent 24.24.24.1
  Return code: Upstream interface index unknown label-switched at stack-depth 1

  Response time 19.05 msec
  Multipath type: IP bitmask
```



```
Address Range 1: 127.0.0.0 ~ 127.0.3.255  
Label Stack:  
Label 1 Value 100000 Protocol LDP
```

```
Hop 22.22.22.4 Depth 3  
Parent 20.20.20.2  
Return code: Egress-ok at stack-depth 1  
Response time 0.79 msec  
Multipath type: None  
Label Stack:  
Label 1 Value 3 Protocol LDP
```


traceroute mpls segment-routing ospf

Syntax

```
traceroute mpls segment-routing ospf<ldp> fec
<destination ip-address>
<detail>
<exp exp>
<fanout fanout-number>
<logical-system logical-system-name>
<no-resolve>
<paths maximum-paths>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source ip-address>
<ttl value>
<update>
<wait seconds>
```

Release Information

Command introduced in Junos OS Release 19.1R1.

Description

Trace route to a remote host for a segment routing label-switched path added by the ISIS protocol. Use **traceroute mpls segment-routing isisp** 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 ip-address—(Optional) Specify the destination address to use when sending probes.

Values: The destination IP address must be within the 127.0.0.0/8 IP address space for Operation, Administration, and Maintenance (OAM) packets.

detail—(Optional) Display detailed output.

exp exp—(Optional) Specify the class-of-service to use when sending probes.

Range: 0 through 7

Default: 7

fanout fanout-number—(Optional) Specify the maximum number of nexthops to search per node.

Range: 1 through 16

Default: 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 *maximum-paths*—(Optional) Specify the maximum number of paths to search.

Range: 1 through 255

Default: 16

retries *retries-number*—(Optional) Specify the number of times to resend probe values.

Range: 1 through 9

Default: 3

routing-instance *routing-instance-name*—(Optional) Specify the name of the routing instance for the 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.

Range: 1 through 125

Default: 64

wait *seconds*—(Optional) Specify the number of seconds to wait before resending a probe.

Range: 5 through 15

Default: 10

Required Privilege Level

network

List of Sample Output

[traceroute mpls segment-routing isis on page 3509](#)

Output Fields

[Table 119 on page 3350](#) describes the output fields for the **traceroute mpls segment-routing isis fec** command and the **traceroute mpls segment-routing isis fec detail** commands. Output fields are listed in the approximate order in which they appear.

Table 142: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls ldp fec command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified

Table 142: traceroute mpls ldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null .	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null .	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail
Response time	Time for the echo request to reach the receiver.	detail
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none .	detail
Label Stack	Label stack used to forward the packet.	detail

Sample Output

traceroute mpls segment-routing isis

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	402006	ISIS	12.1.1.2	(null)	Success
FEC-Stack-Sent: ISIS					
ttl	Label	Protocol	Address	Previous Hop	Probe Status


```

    2  402006  ISIS      23.1.1.2      12.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol  Address        Previous Hop    Probe Status
    3  402006  ISIS      34.1.1.2      23.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol  Address        Previous Hop    Probe Status
    4  402006  ISIS      45.1.1.2      34.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol  Address        Previous Hop    Probe Status
    5           3  ISIS      56.1.1.2      45.1.1.2      Egress
FEC-Stack-Sent: ISIS

```

Path 1 via ge-0/0/0.0 destination 127.0.0.64

```

ttl   Label  Protocol  Address        Previous Hop    Probe Status
    3  402006  ISIS      34.2.1.2      23.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol  Address        Previous Hop    Probe Status
    4  402006  ISIS      45.1.1.2      34.2.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol  Address        Previous Hop    Probe Status
    5           3  ISIS      56.1.1.2      45.1.1.2      Egress
FEC-Stack-Sent: ISIS

```

Path 2 via ge-0/0/0.0 destination 127.0.0.65

```

ttl   Label  Protocol  Address        Previous Hop    Probe Status
    3  402006  ISIS      34.5.1.2      23.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol  Address        Previous Hop    Probe Status
    4  402006  ISIS      45.2.1.2      34.5.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol  Address        Previous Hop    Probe Status
    5           3  ISIS      56.1.1.2      45.2.1.2      Egress
FEC-Stack-Sent: ISIS

```

Path 3 via ge-0/0/0.0 destination 127.0.0.69

traceroute mpls segment-routing isis

Syntax

```
traceroute mpls segment-routing isis<ldp> fec
<destination ip-address>
<detail>
<exp exp>
<fanout fanout-number>
<logical-system logical-system-name>
<no-resolve>
<paths maximum-paths>
<pipe-mode>
<retries retries-number>
<routing-instance routing-instance-name>
<source ip-address>
<ttl value>
<update>
<wait seconds>
```

Release Information

Command introduced in Junos OS Release 19.1R1.

Description

Trace route to a remote host for a segment routing label-switched path added by the ISIS protocol. Use **traceroute mpls segment-routing isisp** 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 ip-address—(Optional) Specify the destination address to use when sending probes.

Values: The destination IP address must be within the 127.0.0.0/8 IP address space for Operation, Administration, and Maintenance (OAM) packets.

detail—(Optional) Display detailed output.

exp exp—(Optional) Specify the class-of-service to use when sending probes.

Range: 0 through 7

Default: 7

fanout fanout-number—(Optional) Specify the maximum number of nexthops to search per node.

Range: 1 through 16

Default: 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 *maximum-paths*—(Optional) Specify the maximum number of paths to search.

Range: 1 through 255

Default: 16

retries *retries-number*—(Optional) Specify the number of times to resend probe values.

Range: 1 through 9

Default: 3

routing-instance *routing-instance-name*—(Optional) Specify the name of the routing instance for the 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.

Range: 1 through 125

Default: 64

wait *seconds*—(Optional) Specify the number of seconds to wait before resending a probe.

Range: 5 through 15

Default: 10

Required Privilege Level

network

List of Sample Output

[traceroute mpls segment-routing isis on page 3513](#)

Output Fields

[Table 119 on page 3350](#) describes the output fields for the **traceroute mpls segment-routing isis fec** command and the **traceroute mpls segment-routing isis fec detail** commands. Output fields are listed in the approximate order in which they appear.

Table 143: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls ldp fec command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified

Table 143: traceroute mpls ldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null .	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null .	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail
Response time	Time for the echo request to reach the receiver.	detail
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none .	detail
Label Stack	Label stack used to forward the packet.	detail

Sample Output

traceroute mpls segment-routing isis

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	402006	ISIS	12.1.1.2	(null)	Success
FEC-Stack-Sent: ISIS					
ttl	Label	Protocol	Address	Previous Hop	Probe Status


```

    2  402006  ISIS      23.1.1.2      12.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol Address      Previous Hop  Probe Status
    3  402006  ISIS      34.1.1.2      23.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol Address      Previous Hop  Probe Status
    4  402006  ISIS      45.1.1.2      34.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol Address      Previous Hop  Probe Status
    5           3  ISIS      56.1.1.2      45.1.1.2      Egress
FEC-Stack-Sent: ISIS

```

Path 1 via ge-0/0/0.0 destination 127.0.0.64

```

ttl   Label  Protocol Address      Previous Hop  Probe Status
    3  402006  ISIS      34.2.1.2      23.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol Address      Previous Hop  Probe Status
    4  402006  ISIS      45.1.1.2      34.2.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol Address      Previous Hop  Probe Status
    5           3  ISIS      56.1.1.2      45.1.1.2      Egress
FEC-Stack-Sent: ISIS

```

Path 2 via ge-0/0/0.0 destination 127.0.0.65

```

ttl   Label  Protocol Address      Previous Hop  Probe Status
    3  402006  ISIS      34.5.1.2      23.1.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol Address      Previous Hop  Probe Status
    4  402006  ISIS      45.2.1.2      34.5.1.2      Success
FEC-Stack-Sent: ISIS
ttl   Label  Protocol Address      Previous Hop  Probe Status
    5           3  ISIS      56.1.1.2      45.2.1.2      Egress
FEC-Stack-Sent: ISIS

```

Path 3 via ge-0/0/0.0 destination 127.0.0.69

CCC and TCC Operational Commands

IN THIS CHAPTER

- [show connections | 3516](#)
- [show route ccc | 3520](#)
- [show route forwarding-table | 3522](#)

show connections

List of Syntax

[Syntax on page 3516](#)

[Syntax \(EX Series Switches\) on page 3516](#)

Syntax

```
show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-interface-switch>
<down | up | up-down>
<history>
<labels>
<logical-system (all | logical-system-name)>
<name>
<status>
```

Syntax (EX Series Switches)

```
show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-interface-switch>
<down | up | up-down>
<history>
<labels>
<name>
<status>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.5 for EX Series switches.

Description

Display information about the configured circuit cross-connect (CCC) connections.

Options

none—Display the standard level of output for all configured CCC connections.

all—(Optional) Display all connections.

brief | extensive—(Optional) Display the specified level of output. Use history to display information about connection history. Use labels to display labels used for transmit and receive LSPs. Use status to display information about the connection and interface status.

- interface-switch**—(Optional) Display interface switch connections only.
- lsp-switch**—(Optional) Display LSP switch connections only.
- p2mp-receive-switch**—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.
- p2mp-transmit-switch**—(Optional) Display local interface to point-to-multipoint LSP switch connections only.
- remote-interface-switch**—(Optional) Display remote interface switch connections only.
- down | up | up-down**—(Optional) Display nonoperational, operational, or both kinds of connections.
- history**—(Optional) Display information about connection history.
- labels**—(Optional) Display labels used for transmit and receive.
- logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.
- name**—(Optional) Display information about the specified connection only.
- status**—(Optional) Display information about the connection and interface status.

Required Privilege Level

view

Output Fields

[Table 83 on page 3145](#) describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 144: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On Off]	Whether link monitoring is enabled: On or Off .
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.

Table 144: show connections Output Fields (*continued*)

Field Name	Field Description
Legend for connection types	Type of connection: <ul style="list-style-type: none"> • if-sw—Layer 2 switching cross-connect. • rmt-if—Remote interface switch. While graceful restart is in progress, rmt-if will display a state (St) of Restart. • lsp-sw—LSP stitching cross-connect. While graceful restart is in progress, lsp-sw will display a state (St) of Restart.
Legend for circuit types	Type of circuits: <ul style="list-style-type: none"> • intf—Interface circuit. • tlsp—Transmit LSP circuit. • rlsp—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the Up (operational) state.
# Up trans	Number of times that the connection or circuit has transitioned to the Up (operational) state.

Sample Output

show connections

user@switch> **show connections**

```
CCC and TCC connections [Link Monitoring On]
Legend for status (St)           Legend for connection types
UN -- uninitialized             if-sw:  interface switching
NP -- not present               rmt-if: remote interface switching
WE -- wrong encapsulation       lsp-sw: LSP switching
DS -- disabled
Dn -- down                      Legend for circuit types
-> -- only outbound conn is up  intf -- interface
```



```
<- -- only inbound conn is up      tlsp -- transmit LSP
Up -- operational                    rlsp -- receive LSP
RmtDn -- remote CCC down
Restart -- restarting
```

CCC Graceful restart : Restarting

Connection/Circuit	Type	St	Time last up	# Up trans
IFSW-ed	if-sw	Up	Aug 5 15:39:15	1
so-1/0/2.0	intf	Up		
t1-0/1/2.0	intf	Up		
SW-db	rmt-if	Restart		0
so-1/0/3.0	intf	Up		
pro4-ca	tlsp	Dn		
pro4-ac	rlsp	NP		

show route ccc

Syntax

```
show route ccc ccc  
<brief | detail | extensive | terse>  
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table.

Options

ccc—Name of an entry with a circuit cross-connect interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

[show connections](#) | [3144](#)

List of Sample Output

[show route ccc extensive on page 3520](#)

Output Fields

For information about output fields, see the output field tables for the *show route* command, the *show route detail* command, the *show route extensive* command, or the *show route terse* command.

Sample Output

show route ccc extensive

```
user@host> show route ccc fe-0/1/0.600 extensive
```



```
mpls.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
fe-0/1/2.600 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel fe-0/1/2.600.0      /16 -> {0.0.0.0}
    *CCC      Preference: 7
              Next-hop reference count: 2
              Next hop: via so-0/0/3.0 weight 0x1, selected
              Label operation: Push 101424
              State: <Active Int>
              Local AS:   100
              Age: 28:13   Metric: 3
              Task: MPLS
              Announcement bits (1): 0-KRT
              AS path: I
```


show route forwarding-table

List of Syntax

[Syntax on page 3522](#)

[Syntax \(MX Series Routers\) on page 3522](#)

[Syntax \(TX Matrix and TX Matrix Plus Routers\) on page 3522](#)

Syntax

```
show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all | vlan-name)>
<vpn vpn>
```

Syntax (MX Series Routers)

```
show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | domain-name)>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<learning-vlan-id learning-vlan-id>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all | vlan-name)>
<vpn vpn>
```

Syntax (TX Matrix and TX Matrix Plus Routers)


```

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<matching matching>
<label name>
<lcc number>
<multicast>
<table routing-instance-name>
<vpn vpn>

```

Release Information

Command introduced before Junos OS Release 7.4.

Option **bridge-domain** introduced in Junos OS Release 7.5

Option **learning-vlan-id** introduced in Junos OS Release 8.4

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: **bridge** (**ccc** | **destination** | **detail** | **extensive** | **interface-name** | **label** | **learning-vlan-id** | **matching** | **multicast** | **summary** | **table** | **vlan** | **vpn**), **ethernet-switching**, **evpn**, **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mcsnoop-inet**, **mcsnoop-inet6**, **mpls**, **satellite-inet**, **satellite-inet6**, **satellite-vpls**, **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 —(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 3530](#)

[show route forwarding-table detail on page 3532](#)

[show route forwarding-table destination extensive \(Weights and Balances\) on page 3533](#)

[show route forwarding-table extensive on page 3533](#)

[show route forwarding-table extensive \(RPF\) on page 3536](#)

[show route forwarding-table extensive \(PIM using point-to-multipoint mode\) on page 3536](#)

[show route forwarding-table \(dynamic list next hop\) on page 3537](#)

[show route forwarding-table family mpls on page 3538](#)

[show route forwarding-table family mpls ccc ge-0/0/1.1004 on page 3538](#)

[show route forwarding-table family vpls on page 3539](#)

[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 3539](#)

[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 3540](#)

[show route forwarding-table family vpls extensive on page 3540](#)

[show route forwarding-table table default on page 3542](#)

[show route forwarding-table table logical-system-name/routing-instance-name on page 3543](#)

[show route forwarding-table vpn on page 3544](#)

Output Fields

[Table 111 on page 3263](#) 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 145: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the table logical-system-name/routing-instance-name option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels

Table 145: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Enabled protocols		All levels

Table 145: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
	<p>The features and protocols that have been enabled for a given routing table. This field can contain the following values:</p> <ul style="list-style-type: none"> • BUM hashing—BUM hashing is enabled. • MAC Stats—Mac Statistics is enabled. • Bridging—Routing instance is a normal layer 2 bridge. • No VLAN—No VLANs are associated with the bridge domain. • All VLANs—The vlan-id all statement has been enabled for this bridge domain. • Single VLAN—Single VLAN ID is associated with the bridge domain. • MAC action drop—New MACs will be dropped when the MAC address limit is reached. • Dual VLAN—Dual VLAN tags are associated with the bridge domain • No local switching—No local switching is enabled for this routing instance.. • Learning disabled—Layer 2 learning is disabled for this routing instance. • MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached. • VPLS—The VPLS protocol is enabled. • No IRB I2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance. • ACKed by all peers—All peers have acknowledged this routing instance. • BUM Pruning—BUM pruning is enabled on the VPLS instance. • Def BD VXLAN—VXLAN is enabled for the default bridge domain. • EVPN—EVPN protocol is enabled for this routing instance. • Def BD OVSDb—Open vSwitch Database (OVSDb) is enabled on the default bridge domain. • Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain. • L2 backhaul—Layer 2 backhaul is enabled. • FRR optimize—Fast reroute optimization • MAC pinning—MAC pinning is enabled for this bridge domain. • MAC Aging Timer—The MAC table aging time is set per routing instance. • EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation. • PBBN—This routing instance is configured as a provider backbone bridged network. 	

Table 145: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • PBN—This routing instance is configured as a provider bridge network. • ETREE—The ETREE protocol is enabled on this EVPN routing instance. • ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance. • Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain. • MPLS control word—Control word is enabled for this MPLS routing instance. 	
Address family	Address family (for example, IP , IPv6 , ISO , MPLS , and VPLS).	All levels
Destination	Destination of the route.	detail extensive
Route Type (Type)	<p>How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses):</p> <ul style="list-style-type: none"> • cloned (clon)—(TCP or multicast only) Cloned route. • destination (dest)—Remote addresses directly reachable through an interface. • destination down (iddn)—Destination route for which the interface is unreachable. • interface cloned (ifcl)—Cloned route for which the interface is unreachable. • route down (ifdn)—Interface route for which the interface is unreachable. • ignore (ignr)—Ignore this route. • interface (intf)—Installed as a result of configuring an interface. • permanent (perm)—Routes installed by the kernel when the routing table is initialized. • user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route Reference (RtRef)	Number of routes to reference.	detail extensive

Table 145: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	<p>Route type flags:</p> <ul style="list-style-type: none"> • none—No flags are enabled. • accounting—Route has accounting enabled. • cached—Cache route. • incoming-iface <i>interface-number</i>—Check against incoming interface. • prefix load balance—Load balancing is enabled for this prefix. • rt nh decoupled—Route has been decoupled from the next hop to the destination. • sent to PFE—Route has been sent to the Packet Forwarding Engine. • static—Static route. 	extensive
Next hop	<p>IP address of the next hop to the destination.</p> <p>NOTE: For static routes that use point-to-point (P2P) outgoing interfaces, the next-hop address is not displayed in the output.</p>	detail extensive
Next hop Type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • discard (dscd)—Discard. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (loc)—Local address on an interface. • routed multicast (mcrt)—Regular multicast next hop. • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp)—Multicast group member. • receive (rcv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail extensive

Table 145: show route forwarding-table Output Fields (continued)

Field Name	Field Description	Level of Output
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	detail extensive none
Next-hop interface (Netif)	Interface used to reach the next hop.	detail extensive none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the Balance field description).	extensive
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	extensive
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when rpf-check is configured on the interface.	extensive

Sample Output

show route forwarding-table

user@host> **show route forwarding-table**

```

Routing table: default.inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif

```



```

default      perm      0      rjct      46      4
0.0.0.0/32   perm      0      dscd      44      1
172.16.1.0/24 ifdn      0      rslv      608      1 ge-2/0/1.0
172.16.1.0/32 iddn      0 172.16.1.0      rcv      606      1 ge-2/0/1.0
172.16.1.1/32 user      0      rjct      46      4
172.16.1.1/32 intf      0 172.16.1.1      locl      607      2
172.16.1.1/32 iddn      0 172.16.1.1      locl      607      2
172.16.1.255/32 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      rcv      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      rcv      610      1 ge-2/0/1.0
10.1.1.1/32  user      0      rjct      46      4
10.1.1.1/32  intf      0 10.1.1.1      locl      611      2
10.1.1.1/32  iddn      0 10.1.1.1      locl      611      2
10.1.1.255/32 iddn      0 ff:ff:ff:ff:ff:ff bcst      609      1 ge-2/0/1.0
10.209.0.0/16 user      0 10.209.63.254      ucst      419      20 fxp0.0
10.209.0.0/16 user      1 0:12:1e:ca:98:0      ucst      419      20 fxp0.0
10.209.0.0/18 intf      0      rslv      418      1 fxp0.0
10.209.0.0/32 dest      0 10.209.0.0      rcv      416      1 fxp0.0
10.209.2.131/32 intf      0 10.209.2.131      locl      417      2
10.209.2.131/32 dest      0 10.209.2.131      locl      417      2
10.209.17.55/32 dest      0 0:30:48:5b:78:d2      ucst      435      1 fxp0.0
10.209.63.42/32 dest      0 0:23:7d:58:92:ca      ucst      434      1 fxp0.0
10.209.63.254/32 dest      0 0:12:1e:ca:98:0      ucst      419      20 fxp0.0
10.209.63.255/32 dest      0 10.209.63.255      bcst      415      1 fxp0.0
10.227.0.0/16 user      0 10.209.63.254      ucst      419      20 fxp0.0

```

...

Routing table: iso

ISO:

```

Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm      0      rjct      27      1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf      0      locl      28      1

```

Routing table: inet6

Internet6:

```

Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm      0      rjct      6      1

```



```

ff00::/8          perm      0          mdsc      4      1
ff02::1/128       perm      0 ff02::1    mcst      3      1

Routing table: ccc
MPLS:
Interface.Label   Type RtRef Next hop          Type Index NhRef Netif
default           perm      0          rjct 16      1
100004(top)fe-0/0/1.0

```

show route forwarding-table detail

user@host> show route forwarding-table detail

```

Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user      2 0:90:69:8e:b1:1b    ucst  132      4 fxp0.0
default          perm      0          rjct   14      1
10.1.1.0/24      intf      0 ff.3.0.21          ucst  322      1 so-5/3/0.0
10.1.1.0/32      dest      0 10.1.1.0           recv  324      1 so-5/3/0.0
10.1.1.1/32      intf      0 10.1.1.1           locl  321      1
10.1.1.255/32    dest      0 10.1.1.255         bcst  323      1 so-5/3/0.0
10.21.21.0/24    intf      0 ff.3.0.21          ucst  326      1 so-5/3/0.0
10.21.21.0/32    dest      0 10.21.21.0         recv  328      1 so-5/3/0.0
10.21.21.1/32    intf      0 10.21.21.1         locl  325      1
10.21.21.255/32  dest      0 10.21.21.255       bcst  327      1 so-5/3/0.0
127.0.0.1/32     intf      0 127.0.0.1          locl  320      1
172.17.28.19/32  clon      1 192.168.4.254       ucst  132      4 fxp0.0
172.17.28.44/32  clon      1 192.168.4.254       ucst  132      4 fxp0.0

...

Routing table: private1__.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm      0          rjct   46      1
10.0.0.0/8       intf      0          rslv  136      1 fxp1.0
10.0.0.0/32      dest      0 10.0.0.0           recv  134      1 fxp1.0
10.0.0.4/32      intf      0 10.0.0.4           locl  135      2
10.0.0.4/32      dest      0 10.0.0.4           locl  135      2

...

Routing table: iso

```



```

ISO:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm    0                rjct   38    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm    0                rjct   22    1
ff00::/8         perm    0                mdsc   21    1
ff02::1/128      perm    0 ff02::1         mcst   17    1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm    0                rjct  28    1

```

show route forwarding-table destination extensive (Weights and Balances)

user@host> **show route forwarding-table destination 3.4.2.1 extensive**

```

Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0
Route interface-index: 0
Flags: sent to PFE
Next-hop type: unilist
Index: 262143 Reference: 1
Nexthop: 172.16.4.4
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
  Flags: sent to PFE
  Nexthop: 00:00:5E:00:53:1b
  Next-hop type: unicast
  Next-hop interface: fxp0.0
                                Route interface-index: 0
                                Index: 132      Reference: 4

```

```

Destination:  default
  Route type:  permanent
  Route reference: 0
  Flags: none
  Next-hop type: reject
                                Route interface-index: 0
                                Index: 14       Reference: 1

```

```

Destination:  127.0.0.1/32
  Route type:  interface
  Route reference: 0
  Flags: sent to PFE
  Nexthop: 127.0.0.1
  Next-hop type: local
                                Route interface-index: 0
                                Index: 320     Reference: 1

```

...

```

Routing table: privatel__inet [Index 1]
Internet:

```

```

Destination:  default
  Route type:  permanent
  Route reference: 0
  Flags: sent to PFE
  Next-hop type: reject
                                Route interface-index: 0
                                Index: 46       Reference: 1

```

```

Destination:  10.0.0.0/8
  Route type:  interface
  Route reference: 0
  Flags: sent to PFE
  Next-hop type: resolve
  Next-hop interface: fxp1.0
                                Route interface-index: 3
                                Index: 136     Reference: 1

```

...

```

Routing table: iso [Index 0]
ISO:

```



```

Destination:  default
  Route type: permanent
  Route reference: 0                      Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                  Index: 38      Reference: 1

```

```

Routing table: inet6 [Index 0]
Internet6:

```

```

Destination:  default
  Route type: permanent
  Route reference: 0                      Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                  Index: 22      Reference: 1

```

```

Destination:  ff00::/8
  Route type: permanent
  Route reference: 0                      Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: multicast discard      Index: 21      Reference: 1

```

```

...

```

```

Routing table: privatel__.inet6 [Index 1]
Internet6:

```

```

Destination:  default
  Route type: permanent
  Route reference: 0                      Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                  Index: 54      Reference: 1

```

```

Destination:  fe80::2a0:a5ff:fe3d:375/128
  Route type: interface
  Route reference: 0                      Route interface-index: 0
  Flags: sent to PFE
  Nexthop: fe80::2a0:a5ff:fe3d:375
  Next-hop type: local                   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 192.0.2.2/30;
    }
  }
}
```

user@host> **show route forwarding-table extensive**

```
Routing table: inet [Index 0]
Internet:
...
...
Destination: 192.0.2.3/32
  Route type: destination
  Route reference: 0                      Route interface-index: 67
  Flags: sent to PFE
  Nexthop: 192.0.2.3
  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 extensive (PIM using point-to-multipoint mode)

user@host> **show route forwarding-table extensive**

```
Destination: 198.51.100.0/24
  Route type: user
  Route reference: 0                      Route interface-index: 335
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: cached, check incoming interface , accounting, sent to PFE, rt nh
decoupled
  Next-hop type: indirect                  Index: 1048575  Reference: 4
  Nexthop:
  Next-hop type: composite                  Index: 627      Reference: 1
```



```

Next-hop type: unicast          Index: 1048574  Reference: 2
Next-hop interface: st0.1, 192.0.2.0

```

show route forwarding-table (dynamic list next hop)

The **show route forwarding table** output shows the two next hop elements for a multihomed EVPN destination.

```
user@host> show route forwarding-table label 299952 extensive
```

```

MPLS:

Destination: 299952
Route type: user
Route reference: 0          Route interface-index: 0
Multicast RPF nh index: 0
P2mpidx: 0
Flags: sent to PFE, rt nh decoupled
Next-hop type: indirect          Index: 1048575  Reference: 2
Nexthop:
Next-hop type: composite          Index: 601      Reference: 2
Next-hop type: indirect          Index: 1048574  Reference: 3
Nexthop: 1.0.0.4
Next-hop type: Push 301632, Push 299776(top) Index: 600 Reference: 2
Load Balance Label: None
Next-hop interface: ge-0/0/1.0
Next-hop type: indirect          Index: 1048577  Reference: 3
Nexthop: 1.0.0.4
Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
Load Balance Label: None
Next-hop interface: ge-0/0/1.0

```

After one of the PE router has been disabled in the EVPN multihomed network, the same **show route forwarding table** output command shows one next hop element and one empty next hop element.

```
user@host> show route forwarding-table label 299952 extensive
```

```

Routing table: default.mpls [Index 0]
MPLS:

Destination: 299952
Route type: user
Route reference: 0          Route interface-index: 0

```



```

Multicast RPF nh index: 0
P2mpidx: 0
Flags: sent to PFE, rt nh decoupled
Next-hop type: indirect          Index: 1048575  Reference: 2
Nexthop:
Next-hop type: composite         Index: 601      Reference: 2
Next-hop type: indirect         Index: 1048577 Reference: 3
Nexthop: 1.0.0.4
Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
Load Balance Label: None
Next-hop interface: ge-0/0/1.0

```

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              rjct   19    1
0                user   0              recv   18    3
1                user   0              recv   18    3
2                user   0              recv   18    3
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 mpls ccc ge-0/0/1.1004

user@host> show route forwarding-table mpls ccc ge-0/0/1.1004

```

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
ge-0/0/1.1004 (CCC) user   0              ulst  1048577 2
                  comp    754      3
                  comp    755      3
                  comp    756      3

Routing table: __mpls-oam__.mpls

```



```

MPLS:
Destination      Type RtRef Next hop      Type Index   NhRef Netif
default          perm    0              dscd    556     1

```

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          dynm    0              flood   353     1
default          perm    0              rjct    298     1
fe-0/1/0.0       dynm    0              flood   355     1
00:00:5E:00:53:1f/48      <<<<<Remote CE

                                dynm    0              indr    351     4
                                Push 800000, Push 100002(top)

so-0/0/0.0
00:00:5E:00:53:1f/48      <<<<<<Local CE

                                dynm    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              dscd    519     1
lsi.1048832      intf    0              indr   1048574     4
                                172.16.3.2    Push 262145     621     2

ge-3/0/0.0
00:00:5E:00:53:01/48 user    0              ucst    590     5 ge-2/3/9.0
0x30003/51       user    0              comp    627     2
ge-2/3/9.0       intf    0              ucst    590     5 ge-2/3/9.0
ge-3/1/3.0       intf    0              ucst    619     4 ge-3/1/3.0
0x30002/51       user    0              comp    600     2
0x30001/51       user    0              comp    597     2

```


show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

user@host> **show route forwarding-table vpls**

```

Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination          Type RtRef Next hop          Type Index   NhRef Netif
default              perm    0                dscd    519     1
lsi.1048834          intf    0                indr 1048574   4
                                172.16.3.2      Push 262145   592    2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user    0                ucst    590     5 ge-2/3/9.0
0x30003/51           user    0                comp    630     2
ge-2/3/9.0           intf    0                ucst    590     5 ge-2/3/9.0
ge-3/1/3.0           intf    0                ucst    591     4 ge-3/1/3.0
0x30002/51           user    0                comp    627     2
0x30001/51           user    0                comp    624     2

```

show route forwarding-table family vpls extensive

user@host> **show route forwarding-table family vpls extensive**

```

Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0                Route interface-index: 72
Flags: sent to PFE
Next-hop type: flood              Index: 289      Reference: 1
Next-hop type: unicast            Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast            Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: default
Route type: permanent
Route reference: 0                Route interface-index: 0
Flags: none
Next-hop type: discard            Index: 341      Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic

```



```

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: 00:00:5E:00:53:05/48
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                rjct   36    2
0.0.0.0/32       perm    0                dscd   34    1
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	
172.16.0.1/32	user	0		dscd	561	2	
172.16.2.0/24	intf	0		rslv	771	1	ge-1/2/0.3
172.16.2.0/32	dest	0	172.16.2.0	recv	769	1	ge-1/2/0.3
172.16.2.1/32	intf	0	172.16.2.1	locl	770	2	
172.16.2.1/32	dest	0	172.16.2.1	locl	770	2	
172.16.2.2/32	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
172.16.2.255/32	dest	0	172.16.2.255	bcst	768	1	ge-1/2/0.3


```

172.16.233.0/4      perm      1      mdsc    562      1
172.16.233.1/32    perm      0 172.16.233.1    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
172.16.233.0/4       perm      2      mdsc      5      3
172.16.233.1/32     perm      0 172.16.233.1    mcst      1      8

```


172.16.233.5/32	user	1	172.16.233.5	mcst	1	8
255.255.255.255/32	perm	0		bcst	2	3

On QFX5200, the results for this command look like this:

show route forwarding-table family mpls

```

Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 65 1
0 user 0 recv 64 4
1 user 0 recv 64 4
2 user 0 recv 64 4
13 user 0 recv 64 4
300384 user 0 9.1.1.1 Pop 1711 2 xe-0/0/34.0
300384(S=0) user 0 9.1.1.1 Pop 1712 2 xe-0/0/34.0
300400 user 0 ulst 131071 2
                                10.1.1.2 Pop 1713 1 xe-0/0/38.0
                                172.16.11.2 Pop 1714 1 xe-0/0/40.0
300400(S=0) user 0 ulst 131072 2
                                10.1.1.2 Pop 1715 1 xe-0/0/38.0
                                172.16.11.2 Pop 1716 1 xe-0/0/40.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 1681 1

```


PCEP Operational Commands

IN THIS CHAPTER

- [clear path-computation-client statistics | 3548](#)
- [request path-computation-client active-pce | 3550](#)
- [show isis spring sensor info | 3551](#)
- [show path-computation-client active-pce | 3554](#)
- [show path-computation-client lsp | 3559](#)
- [show path-computation-client statistics | 3565](#)
- [show path-computation-client status | 3573](#)
- [show path-computation-client traffic-steering | 3576](#)
- [show spring-traffic-engineering | 3578](#)

clear path-computation-client statistics

Syntax

```
clear path-computation-client statistics  
<pce-id | all>
```

Release Information

Statement introduced in Junos OS Release 12.3.

Statement introduced in Junos OS Release 16.1R3 for QFX Series switches.

Command introduced in Junos OS Release 17.1R1 for ACX Series routers.

Description

Clear Path Computation Element (PCE) statistics.

Options

pce-id—(Optional) Clear statistics of the specified PCE.

all—(Optional) Clear statistics of all available PCEs configured on the path computation client (PCC).

Required Privilege Level

clear

RELATED DOCUMENTATION

[show path-computation-client statistics](#) | [3565](#)

List of Sample Output

[clear path-computation-client statistics pce-id on page 3548](#)

[clear path-computation-client statistics all on page 3549](#)

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

```
clear path-computation-client statistics pce-id
```

```
user@host> clear path-computation-client statistics pce1
```


clear path-computation-client statistics all

user@host> clear path-computation-client statistics all

request path-computation-client active-pce

Syntax

```
request path-computation-client active-pce pce-id
```

Release Information

Command introduced in Junos OS Release 12.3.

Command introduced in Junos OS Release 16.1R3 for QFX Series switches.

Command introduced in Junos OS Release 17.1R1 for ACX Series routers.

Description

Request a new active Path Computation Element (PCE).

Options

pce-id—Unique user defined ID for this PCE.

retry-delegation—Retry label-switched path (LSP) delegation

Required Privilege Level

request

RELATED DOCUMENTATION

[show path-computation-client active-pce](#) | [3554](#)

List of Sample Output

[request path-computation-client active-pce pce-id](#) on page 3550

Sample Output

```
request path-computation-client active-pce pce-id
```

```
user@host> request path-computation-client active-pce pce1
```


show isis spring sensor info

Syntax

```
show isis spring sensor info
logical-system (all | logical-system-name)
```

Release Information

Command introduced in 19.1R1 on MX Series routers with MPC and MIC interfaces, and PTX series routers.

Description

Displays a list of sensors associated with the label IS-IS route and next hops for segment routing traffic. The command only displays the information related to the sensors and not the traffic statistics.

Options

none— Display the sensor information of an IS-IS SPRING route.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

<i>sensor-based-stats</i>
<i>source-packet-routing (Protocols IS-IS and OSPF)</i>
<i>Understanding Source Packet Routing in Networking (SPRING)</i>

List of Sample Output

[show isis spring sensor info on page 3552](#)

Output Fields

Table 1 describes the output fields for the **show isis spring sensor info** command. Output fields are listed in the approximate order in which they appear.

Table 146: show isis spring sensor info Output Fields

Field	Field Description
Sensor-name	Represents the router or interface that the sensor is associated with.

Table 146: show isis spring sensor info Output Fields (*continued*)

Field	Field Description
Sensor-id	Unique number associated either with route or interface.

Sample Output

show isis spring sensor info

user@host> **show isis spring sensor info**

Per-interface-per-member-link Ingress Sensor:

```
-----
Sensor-name                Sensor-id
aggr_ingress_intf_sensor   3221225484
```

Per-interface-per-member-link Egress Sensor:

```
-----
Sensor-name                Sensor-id
ge-0/0/0.0                 3221225497
ge-0/0/1.0                 3221225498
ge-0/0/2.0                 3221225499
```

Per-sid Ingress Sensor:

```
-----
Sensor-name                Sensor-id
16                          3221225478
17                          3221225479
18                          3221225474
19                          3221225475
20                          3221225482
21                          3221225483
22                          3221225480
23                          3221225481
24                          3221225489
25                          3221225490
400001                     3221225491
400002                     3221225492
400005                     3221225487
400006                     3221225488
400009                     3221225493
400010                     3221225494
```


400011	3221225495
400012	3221225496
IPv4/IPv6 Per-sid Egress Sensor:	

Sensor-name	Sensor-id
L-ISIS-::10.10.10.1	3221225474

show path-computation-client active-pce

Syntax

```
show path-computation-client active-pce
<brief | detail>
```

Release Information

Command introduced in Junos OS Release 12.3.
 Command introduced in Junos OS Release 16.1R3 for QFX Series switches.
 Command introduced in Junos OS Release 17.1R1 for ACX Series routers.

Description

Displays information about the current active Path Computation Element (PCE).

Options

- none**—Display brief information about the current active PCE.
- brief | detail**—(Optional) Display the specific level of output.

Required Privilege Level

view

RELATED DOCUMENTATION

| [request path-computation-client active-pce](#) | [3550](#)

List of Sample Output

- [show path-computation-client active-pce on page 3557](#)
- [show path-computation-client active-pce detail on page 3557](#)

Output Fields

[Table 147 on page 3554](#) describes the output fields for the **show path-computation-client active-pce** command. Output fields are listed in the approximate order in which they appear.

Table 147: show path-computation-client active-pce Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the current active PCE.	All levels
Priority	Active PCE priority.	All levels

Table 147: show path-computation-client active-pce Output Fields (*continued*)

Field Name	Field Description	Level of Output
PCE status	<p>Active PCE state:</p> <ul style="list-style-type: none"> • PCE_STATE_NEW—Initial PCEP session state. • PCE_STATE_RECONNECT—Trying to re-establish TCP connection with the PCEP peer. • PCE_STATE_CONNECTING—Establishing TCP connection with the PCEP peer. • PCE_STATE_CONNECTED—TCP connection established with the PCEP peer. • PCE_STATE_SYNC—Open messages exchanged with the PCEP peer and entering SYNC state. • PCE_STATE_UP—PCEP session established. 	All levels
Session type	<p>Active PCE type:</p> <ul style="list-style-type: none"> • PCE_TYPE_STATELESS—Does not learn LSP state information from PCC. • PCE_TYPE_STATEFUL—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update LSP state. A PCC maintains synchronization with the PCE. • PCE_TYPE_STATEFULACTIVE—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegate control of their LSPs to the PCE. 	All levels
PCE-mastership	<p>PCE mastership state:</p> <ul style="list-style-type: none"> • main—Current active PCE. • backup—Backup PCE. 	All levels
PCRpts	Number of PC report (PCRpt) messages sent by PCC to a stateful PCE to report current state of LSP(s).	All levels
PCUpdates	Number of PC update (PCUpd) messages sent by a PCE to a PCC to update LSP parameters.	All levels
Local Keepalive timer	Keepalive timer used by or for the PCC.	All levels
Local Dead timer	Dead timer used by or for the PCC.	All levels
Remote Keepalive timer	Keepalive timer used by or for the PCE.	All levels

Table 147: show path-computation-client active-pce Output Fields (*continued*)

Field Name	Field Description	Level of Output
Remote Dead timer	Dead timer used by or for the PCE.	All levels
PCErr-recv	Information about type, value, and number of PC Error messages received.	All levels
Max unknown messages	Maximum number of unknown messages received for a PCEP session. Recommended value is 5. If the number of unknown messages received by a PCC or PCE is greater than or equal to the maximum number, the PCEP session is closed.	detail
Keepalives received	Number of Keepalive messages received by a PCC from a PCE.	detail
Keepalives sent	Number of Keepalive messages sent by a PCC to a PCE.	detail
Dead timer	Dead timer used by the current active PCE.	detail
Elapsed as main current	Time (in seconds) the PCE is in the main mastership state.	detail
Elapsed as main total	Time (in seconds) the PCE became main from the last PCCD restart.	detail
Unknown msgs/min rate	Number of unknown messages received per minute.	detail
Session failures	Number of PCEP session failures with the PCE.	detail
Delegation timeout in	Time (in seconds) left for LSP delegation to timeout.	detail
Delegation failures	Number of LSP delegation failures.	detail
Connection down	Time (in seconds) since the PCEP session is down.	detail
PCErr-sent	Information about type, value, and number of PC Error messages sent.	All levels

Sample Output

show path-computation-client active-pce

```
user@host> show path-computation-client active-pce
```

```
PCE pce1
General
  IP address           : 10.209.57.166
  Priority              : 2
  PCE status            : PCE_STATE_NEW
  Session type          : PCE_TYPE_STATEFULACTIVE
  PCE-mastership        : main

Counters
  PCReqs                Total: 0          last 5min: 0          last hour: 0
  PCReps                Total: 0          last 5min: 0          last hour: 0
  PCRpts                Total: 0          last 5min: 0          last hour: 0
  PCUpdates             Total: 0          last 5min: 0          last hour: 0

Timers
  Local                Keepalive timer:    0 [s]   Dead timer:    0 [s]
  Remote               Keepalive timer:    0 [s]   Dead timer:    0 [s]

Errors
  PCErr-recv
  PCErr-sent
    Type: 19          Value: 3          Count: 1
  PCE-PCC-NTFS
  PCC-PCE-NTFS
```

show path-computation-client active-pce detail

```
user@host> show path-computation-client active-pce detail
```

```
PCE pce1
General
  IP address           : 172.22.25.223
  Priority              : 1
```



```

PCE status           : PCE_STATE_RECONNECT
Session type         : PCE_TYPE_STATEFULACTIVE
PCE-mastership       : main
Max unknown messages : 5
Keepalives received  : 0
Keepalives sent       : 0
Dead timer           : 0 [s]
Elapsed as main current : 1 [s]
Elapsed as main total : 2542 [s]
Unknown msgs/min rate : 0
Session failures     : 575
Delegation timeout in : 14 [s]
Delegation failures   : 21928
Connection down       : 16 [s]

```

Counters

PCReqs	Total: 0	last 5min: 0	last hour: 0
PCReps	Total: 0	last 5min: 0	last hour: 0
PCRpts	Total: 31512	last 5min: 7243	last hour:
7243			
PCUpdates	Total: 80	last 5min: 40	last hour: 40

Timers

Local	Keepalive timer:	30 [s]	Dead timer:	120 [s]
Remote	Keepalive timer:	30 [s]	Dead timer:	120 [s]

Errors

```

PCErr-recv
PCErr-sent
Type: 1      Value: 2      Count: 12
PCE-PCC-NTFS
PCC-PCE-NTFS

```


show path-computation-client lsp

Syntax

```
show path-computation-client lsp
<extensive>
<p2mp>
```

Release Information

Command introduced in Junos OS Release 17.2R1 on MX Series routers.

extensive and **p2mp** options introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display the state of label-switched paths (LSPs) known to the Path Computation Client (PCC).

Options

none—Display information about LSPs known to the PCC.

extensive—(Optional) Display extensive level of output about each known LSP - point-to-point and point-to-multipoint LSPs.

p2mp—(Optional) Display information about known point-to-multipoint Path Computation Element (PCE)-initiated LSPs.

Required Privilege Level

view

RELATED DOCUMENTATION

[show path-computation-client active-pce | 3554](#)

[show path-computation-client statistics | 3565](#)

[show path-computation-client status | 3573](#)

List of Sample Output

[show path-computation-client lsp on page 3562](#)

[show path-computation-client lsp p2mp on page 3562](#)

[show path-computation-client lsp extensive \(PCE-Initiated Point-to-Multipoint LSP Mapping to MVPN\) on page 3563](#)

Output Fields

[Table 148 on page 3560](#) describes the output fields for the **show path-computation-client lsp** command. Output fields are listed in the approximate order in which they appear.

Table 148: show path-computation-client lsp Output Fields

Field Name	Field Description
LSP Name	Name of the LSP.
Status	LSP status: <ul style="list-style-type: none"> • Primary(Act)—Primary LSP. • Bypass—PCE-initiated bypass LSP.
PLSP-Id	PCEP-specific unique identifier for each LSP. The ID is created by the PCC for the lifetime of a PCEP session.
LSP-Type	Type of LSP: <ul style="list-style-type: none"> • External provisioned • Local
Controller	Name of the external path computing entity.
Path-Setup-Type	Protocol used to set up the LSP: <ul style="list-style-type: none"> • RSVP-TE • SPRING-TE
Template	Name of template used.
P2MP name	Name of the point-to-multipoint tree that includes the sub-LSPs of the PCE-initiated LSP.
P2MP Branch Name	Name of the branch sub-LSP that makes up the point-to-multipoint tree of the PCE-initiated LSP.
PathName	Name of LSP path.
From	Ingress IP address of LSP.
To	Egress IP address of LSP.
State	LSP state: up, down.
Active Path	Name of active path.
Link Protection	LSP Link protection.

Table 148: show path-computation-client lsp Output Fields (*continued*)

Field Name	Field Description
P2mp tree	Name of the point-to-multipoint tree that includes the sub-LSPs of the PCE-initiated LSP.
Path cspf status	CSPF computation done by an external controller or local router.
LSP-ID	LSP identifier.
RSVP Error	RSVP error ID.
Priorities	Setup and hold priorities.
Bandwidth	LSP bandwidth.
Requested AutoBw	Requested bandwidth to controller for auto-bandwidth.
Controller	Name of external controller.
Record Route	LSP record route object.
From PCE ERO (received)	Explicit Route Object (ERO) received by controller from the routing protocol process.
From RPD ERO (reported)	ERO reported from the routing protocol process.
Configured ERO on PCC	ERO configured on router.
PCE Traffic Steering	MVPN flow specification capability.
FS-ID	Flow specification ID.
Route Distinguisher	Route Distinguisher of MVPN instance.
Source Prefix	Matching source address of MVPN flow specification.
Multicast Group Prefix	Matching group address of MVPN flow specification.
State	Flow specification state of specific flow specification ID.
Last Rpt/Pcreqest received from RPD at	Time the last PC report or PE request message was received.
Last Update sent to PCE at	Time the last update was sent to PCE.

Table 148: show path-computation-client lsp Output Fields (*continued*)

Field Name	Field Description
Last PcUpdate/PcCreate received from PCE at	Time the last PC update message was received from the controller.
Last error sent to PCE at	Time the last PC error was sent to the controller.
Last 5 reasons to send Report/Pcrequest	<ul style="list-style-type: none"> • Reconfig • Down • Get_info • Up • Active

Sample Output

show path-computation-client lsp

```
user@host> show path-computation-client lsp
```

Name	Status	PLSP-Id	LSP-Type	Controller	
Path-Setup-Type	Template	p2mp-name			
lsp1	Primary(Act)	1	ext-provised	pce1	rsvp
	default_pvc	-			
lsp1	Bypass	2	ext-provised	pce1	rsvp
	default_pvc	-			

show path-computation-client lsp p2mp

```
user@host> show path-computation-client lsp p2mp
```

```

P2MP name: p2mp_tree1

P2MP Branch Name: p2mp_tree1_leaf1

P2MP Branch Name: p2mp_tree1_leaf2

P2MP name: p2mp_tree2

P2MP Branch Name: p2mp_tree2_leaf1

```


P2MP Branch Name: p2mp_tree2_leaf2

show path-computation-client lsp extensive (PCE-Initiated Point-to-Multipoint LSP Mapping to MVPN)

user@host> show path-computation-client lsp extensive

```

LSP Name           : lsp_test
PathName           : primary
From : 128.220.10.163 To : 128.220.10.151 State : Up
Active Path        : primary
Link Protection none
LSP Type           : ext-provised
P2mp tree          : p2mp_test
Path cspf status    : external_cspf
Template           : default_pvc
PLSP-ID            : 2
LSP-ID             : 1
RSVP Error          : 0x0
Priorities          : 0
Bandwidth           : 98760
Requested AutoBw    : 0
Controller          : pcel
Record Route        : 1.4.0.2(S) 1.1.0.4(S) 4.7.0.2(S) 1.1.0.4(S) 7.8.0.2(S)
                     1.1.0.0(S)
From PCE ERO (received) : 128.220.10.15(L)
From RPD ERO (reported) : 128.220.10.15(L)
Configured ERO on PCC   : 128.220.10.15(L)
PCE Traffic Steering   :
  FS-ID: 1
    Route Distinguisher : 12345112:123
    Source Prefix.       : 10.220.10.15/24
    Multicast Group Prefix : 10.220.10.11/24
    State                : Active(MVPN instance configured)
  FS-ID: 2
    Route Distinguisher. : 12345112:123
    Source Prefix.       : 10.220.10.15/24
    Multicast Group Prefix : 10.220.10.12/24
    State                : Inactive(MVPN S,G already exist)
Last Rpt/Pcrequest received from RPD at : 23:46:12.000
Last Update sent to PCE at              : 05:30:00.000
Last PcUpdate/PcCreate received from PCE at : 23:46:12.000
Last error sent to PCE at               : 05:30:00.000

```


Last 5 reasons to send Report/Pcrequest : Reconfig, Down, Get_info, Up,
Active,

show path-computation-client statistics

Syntax

```
show path-computation-client statistics  
<brief | detail>  
<all>
```

Release Information

Command introduced in Junos OS Release 12.3.

Command introduced in Junos OS Release 16.1R3 for QFX Series switches.

Command introduced in Junos OS Release 17.1R1 for ACX Series routers.

Description

Display statistics about the Path Computation Element (PCE).

Options

none—Display statistics about the primary PCE.

brief | detail—(Optional) Display the specific level of output.

all—(Optional) Display the statistics about all PCEs configured on the PCC.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear path-computation-client statistics](#) | 3548

List of Sample Output

[show path-computation-client statistics all on page 3568](#)

[show path-computation-client statistics detail on page 3570](#)

[show path-computation-client statistics \(PCE-Initiated Point-to-Multipoint LSP Mapping to MVPN\) on page 3571](#)

Output Fields

[Table 149 on page 3566](#) describes the output fields for the **show path-computation-client statistics** command. Output fields are listed in the approximate order in which they appear.

Table 149: show path-computation-client statistics Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the PCE.	All levels
Priority	PCE priority.	All levels
PCE status	PCE state: <ul style="list-style-type: none"> • PCE_STATE_NEW—Initial PCEP session state. • PCE_STATE_RECONNECT—Trying to re-establish TCP connection with the PCEP peer. • PCE_STATE_CONNECTING—Establishing TCP connection with the PCEP peer. • PCE_STATE_CONNECTED—TCP connection established with the PCEP peer. • PCE_STATE_SYNC—Open messages exchanged with the PCEP peer and entering SYNC state. • PCE_STATE_UP—PCEP session established. 	All levels
Session type	Active PCE type: <ul style="list-style-type: none"> • PCE_TYPE_STATELESS—Does not learn LSP state information from PCC. • PCE_TYPE_STATEFUL—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update LSP state. A PCC maintains synchronization with the PCE. • PCE_TYPE_STATEFULACTIVE—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegate control of their LSPs to the PCE. 	All levels
PCE-mastership	PCE mastership state: <ul style="list-style-type: none"> • main • primary • backup 	All levels

Table 149: show path-computation-client statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
PCRpts	Number of PC report (PCRpt) messages sent by PCC to a stateful PCE to report current state of LSP(s).	All levels
PCUpdates	Number of PC update (PCUpd) messages sent by a PCE to a PCC to update LSP parameters.	All levels
Local Keepalive timer	Keepalive timer used by or for the PCC.	All levels
Local Dead timer	Dead timer used by or for the PCC.	All levels
Remote Keepalive timer	Keepalive timer used by or for the PCE.	All levels
Remote Dead timer	Dead timer used by or for the PCE.	All levels
PCErr-recv	Information about type, value, and number of PC Error messages received.	All levels
PCErr-sent	Information about type, value, and number of PC Error messages sent.	All levels
Max unknown messages	Maximum number of unknown messages received for a PCEP session. Recommended value is 5. If the number of unknown messages received by a PCC or PCE is greater than or equal to the maximum number, the PCEP session is closed.	detail
Keepalives received	Number of Keepalive messages received by a PCC from a PCE.	detail
Keepalives sent	Number of Keepalive messages sent by a PCC to a PCE.	detail
Elapsed as main current	Time (in seconds) the PCE is in the main mastership state.	detail
Elapsed as main total	Time (in seconds) the PCE became main from the last PCCD restart.	detail

Table 149: show path-computation-client statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Unknown msgs/min rate	Number of unknown messages received per minute.	detail
Session failures	Number of PCEP session failures with the PCE.	detail
Delegation timeout in	Time (in seconds) left for LSP delegation to timeout.	detail
Delegation failures	Number of LSP delegation failures.	detail
Connection down	Time (in seconds) since the PCEP session is down.	detail
Local IP address	IP address of the PCC.	detail
LSP provisioning allowed	LSP provisioning capability.	detail
P2MP LSP report allowed	Report capability of point-to-multipoint LSP.	detail
P2MP LSP update allowed	Update capability of point-to-multipoint LSP.	detail
P2MP LSP init allowed	Initiate capability of point-to-multipoint LSP.	detail
PCE Traffic Steering	Traffic steering capability of the PCE.	detail

Sample Output

show path-computation-client statistics all

user@host> **show path-computation-client statistics all**

```
PCE pce1
```


General

IP address : 10.209.57.166
 Priority : 2
 PCE status : PCE_STATE_NEW
 Session type : PCE_TYPE_STATEFULACTIVE
 PCE-mastership : main

Counters

PCReqs	Total: 0	last 5min: 0	last hour: 0
PCReps	Total: 0	last 5min: 0	last hour: 0
PCRpts	Total: 0	last 5min: 0	last hour: 0
PCUpdates	Total: 0	last 5min: 0	last hour: 0

Timers

Local	Keepalive timer:	0 [s]	Dead timer:	0 [s]
Remote	Keepalive timer:	0 [s]	Dead timer:	0 [s]

Errors

PCErr-recv
 PCErr-sent
 PCE-PCC-NTFS
 PCC-PCE-NTFS

PCE pce2

General

IP address : 10.31.32.1
 Priority : 10
 PCE status : PCE_STATE_NEW
 Session type : PCE_TYPE_STATEFULACTIVE
 PCE-mastership : backup

Counters

PCReqs	Total: 0	last 5min: 0	last hour: 0
PCReps	Total: 0	last 5min: 0	last hour: 0
PCRpts	Total: 0	last 5min: 0	last hour: 0
PCUpdates	Total: 0	last 5min: 0	last hour: 0


```

Timers
  Local          Keepalive timer:      0 [s]  Dead timer:      0 [s]
  Remote         Keepalive timer:      0 [s]  Dead timer:      0 [s]

Errors
  PCErr-recv
  PCErr-sent
  PCE-PCC-NTFS
  PCC-PCE-NTFS

```

show path-computation-client statistics detail

```
user@host> show path-computation-client statistics detail
```

```

PCE pcel
General
  IP address          : 10.209.57.166
  Priority             : 2
  PCE status          : PCE_STATE_NEW
  Session type        : PCE_TYPE_STATEFULACTIVE
  PCE-mastership      : main
  Max unknown messages : 5
  Keepalives received : 0
  Keepalives sent     : 0
  Dead timer          : 0 [s]
  Elapsed as main current : 294 [s]
  Elapsed as main total : 294 [s]
  Unknown msgs/min rate : 0
  Session failures    : 0
  Replies timedout    : 0
  Delegation timeout in : 26 [s]
  Delegation failures  : 0
  Connection down     : 4 [s]

Counters
  PCReqs          Total: 0          last 5min: 0          last hour: 0
  PCReps          Total: 0          last 5min: 0          last hour: 0
  PCRpts          Total: 0          last 5min: 0          last hour: 0
  PCUpdates       Total: 0          last 5min: 0          last hour: 0

```



```

Timers
  Local      Keepalive timer:      0 [s]  Dead timer:      0 [s]
  Remote     Keepalive timer:      0 [s]  Dead timer:      0 [s]

Errors
  PCErr-recv
  PCErr-sent
  PCE-PCC-NTFS
  PCC-PCE-NTFS

```

show path-computation-client statistics (PCE-Initiated Point-to-Multipoint LSP Mapping to MVPN)

user@host> show path-computation-client statistics

```

PCE pcel
-----
General
  PCE IP address      : 10.220.11.59
  Local IP address    : 128.220.11.56
  Priority             : 0
  PCE status          : PCE_STATE_UP
  Session type        : PCE_TYPE_STATEFULACTIVE
  LSP provisioning allowed : On
  P2MP LSP report allowed : On
  P2MP LSP update allowed : On
  P2MP LSP init allowed  : Off
  PCE-mastership       : main
  PCE Traffic Steering  : On

Counters
  PCReqs      Total: 0      last 5min: 0      last hour: 0
  PCReps      Total: 0      last 5min: 0      last hour: 0
  PCRpts      Total: 4      last 5min: 0      last hour: 0
  PCUpdates   Total: 2      last 5min: 0      last hour: 0
  PCCreates   Total: 1      last 5min: 0      last hour: 0

Timers
  Local  Keepalive timer:  30 [s]  Dead timer:  120 [s]  LSP cleanup timer:
500 [s]
  Remote Keepalive timer:  30 [s]  Dead timer:  120 [s]  LSP cleanup timer:
0 [s]

```


Errors

PCErr-recv

PCErr-sent

PCE-PCC-NTFS

PCC-PCE-NTFS

Pcupdate empty ero action counters

Send-err : 0

Tear down path : 0

Routing decision : 0

Routing decision failed: 0

show path-computation-client status

Syntax

```
show path-computation-client status  
<extensive>
```

Release Information

Command introduced in Junos OS Release 17.2R1 on MX Series routers.

extensive option introduced in Junos OS Release 18.3R1 on MX Series routers.

Description

Display the status of the Path Computation Client (PCC).

Options

none—Display the status of the PCC.

extensive—(Optional) Display extensive information about the PCC including point-to-point and point-to-multipoint PCE-initiated LSPs.

For point-to-multipoint PCE-initiated LSPs, the extensive output displays the point-to-multipoint LSP tree and branches separately for a PCEP session.

Required Privilege Level

view

RELATED DOCUMENTATION

[show path-computation-client active-pce | 3554](#)

[show path-computation-client statistics | 3565](#)

[show path-computation-client lsp | 3559](#)

List of Sample Output

[show path-computation-client status on page 3574](#)

Output Fields

[Table 150 on page 3574](#) describes the output fields for the **show path-computation-client status** command. Output fields are listed in the approximate order in which they appear.

Table 150: show path-computation-client status Output Fields

Field Name	Field Description
Session	Name of the PCE with which the PCEP session is established.
Type	Type of PCE.
Provisioning	Provisioning status of the PCE.
Status	PCEP session status.
Total number of LSPs	Number of LSPs in total.
Static LSPs	Status of point-to-point and point-to-multipoint static LSPs.
Externally controlled LSPs	Status of point-to-point and point-to-multipoint LSPs that are controlled by a PCE.
Externally provisioned LSPs	Status of point-to-point and point-to-multipoint LSPs that are provisioned by a PCE.
Orphaned LSPs	Status of LSPs that are in the orphaned state because of PCEP session failure.
Delegated	Status of point-to-point and point-to-multipoint LSPs that are delegated to the PCE by the PCC.

Sample Output

show path-computation-client status

user@host> **show path-computation-client status extensive**

Session	Type	Provisioning	Status
pce1	Stateful Active	On	Up
LSP Summary			
Total number of LSPs			: 0
Static LSPs			: 0
P2P			: 0
			: 0/0 (primary/bypass)
P2MP			: 0/0 (branches/trees)
Externally controlled LSPs			: 0
P2P			: 0


```

                : 0/0 (primary/bypass)
      P2MP                : 0/0 (branches/trees)
Externally provisioned LSPs : 0/16000 (current/limit)
      P2P                : 0
                : 0/0 (primary/ bypass)
      P2MP                : 0/0 (branches/trees)
Orphaned LSPs            : 0

pcel (main)
  Delegated                : 0
    P2P                    : 0
                : 0/0 (primary/bypass)
    P2MP                    : 0/0 (branches/trees)
Externally provisioned: 0
    P2P                    : 0
                : 0/0 (primary/bypass)
    P2MP                    : 0/0 (branches/trees)

```


show path-computation-client traffic-steering

Syntax

```
show path-computation-client traffic-steering
pce pce-id
```

Release Information

Command introduced in Junos OS Release 19.4R1 on all platforms.

Description

Display information on flow specification capability (also called traffic steering functionality). This is used for enabling the mapping of PCE-initiated point-to-multipoint LSPs to an MVPN routing-instance.

Options

none—Display information about the flow specification capability.

pce pce-id—(Optional) Display information about the flow specification capability for the specified PCE ID.

Required Privilege Level

view

RELATED DOCUMENTATION

| [show path-computation-client statistics](#) | [3565](#)

List of Sample Output

[show path-computation-client traffic-steering on page 3577](#)

Output Fields

[Table 151 on page 3576](#) describes the output fields for the **show path-computation-client traffic-steering** command. Output fields are listed in the approximate order in which they appear.

Table 151: show path-computation-client traffic-steering Output Fields

Field Name	Field Description
FS-ID	
Route Distinguisher	
Source Prefix	

Table 151: show path-computation-client traffic-steering Output Fields (continued)

Field Name	Field Description
Multicast Group Prefix	
State	
PCE Traffic Steering	

Sample Output

show path-computation-client traffic-steering

```
user@host> show path-computation-client traffic-steering
```

```
PCE Traffic Steering: pce1
-----
FS-ID: 1
  P2MP LSP           : p2mp_lsp1
  Route Distinguisher : 12345112:123
  Source Prefix       : 10.220.10.15/24
  Multicast Group Prefix : 10.220.10.11/24
  State               : Active(MVPN instance configured)
FS-ID: 2
  P2MP LSP           : p2mp_lsp1
  Route Distinguisher : 12345112:123
  Source Prefix       : 10.220.11.16/24
  Multicast Group Prefix : 10.220.10.11/24
  State               : Inactive(MVPN S,G already exist)
PCE Traffic Steering: pce2
-----
FS-ID: 1
  P2MP LSP           : p2mp_lsp2
  Route Distinguisher : 12345333:321
  Source Prefix       : 10.222.21.16/24
  Multicast Group Prefix : 10.222.11.11/24
  State               : Active(MVPN instance configured)
```


show spring-traffic-engineering

Syntax

```
show spring-traffic-engineering (lsp | overview | sbfd)
<brief | detail>
<logical-system (all | logical-system-name)>
<name lsp-name>
```

Release Information

Command introduced in Junos OS Release 17.2 on MX Series routers.

sbfd option introduced in Junos OS Release 19.4R1 on all platforms.

Description

Display ingress details of SPRING traffic engineering.

Options

brief | detail—(Optional) Display the specific level of output.

lsp—Display details of SPRING traffic engineered LSPs on the ingress router or the Path Computation Client (PCC).

overview—Display overview of SPRING traffic engineered LSPs on the ingress router, or the PCC.

sbfd—Display SPRING Traffic Engineered BFD session.

name *lsp-name*—(Optional) Regular expression for LSP names to match for displaying SPRING traffic engineering details.

Required Privilege Level

view

RELATED DOCUMENTATION

[Support of SPRING-TE for the Path Computation Element Protocol Overview | 1777](#)

[Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs | 1782](#)

List of Sample Output

[show spring-traffic-engineering lsp name on page 3580](#)

[show spring-traffic-engineering lsp detail on page 3580](#)

[show spring-traffic-engineering overview on page 3580](#)

[show spring-traffic-engineering sbfd detail on page 3580](#)

[show spring-traffic-engineering sbfd name <name> detail on page 3581](#)

Output Fields

Table 152 on page 3579 describes the output fields for the **show spring-traffic-engineering** command. Output fields are listed in the approximate order in which they appear.

Table 152: show spring-traffic-engineering Output Fields

Field Name	Field Description
To	IP address of the SPRING-TE LSP destination.
State	State of the SPRING-TE LSP: <ul style="list-style-type: none"> • Up • Down
LSP Name	Name of the SPRING-TE LSP.
S-ERO	Source Explicit Route Object (ERO), or LSP path.
Bandwidth	Bandwidth allocated for the SPRING-TE LSP.
Route preference	Route preference of the SPRING-TE LSP.
Number of LSPs	Statistics of the total number of SPRING-TE LSPs and the LSP state.
External controllers	Name of the LSP external controller. By default the only supported external controller is pccd .
BFD name	Name of the BFD session. The name is auto-generated in the V4-srte_bfd_session-id for IPv6. The name is based on the Explicit Route Object (ERO) stack of the LSP path, that is, if multiple LSPs have same path they share the same BFD session name.
BFD status	Status of the BFD session: UP, DOWN.
Referencing LSPs	Name of referencing LSP. If the LSP does not have a path name, then the referencing LSP is displayed as unnamed path .
SR-ERO hop count	Number of hops in the segment routing ERO.
Hop 1	Represents the path of the BFD session. If any other LSP is on same path, it has the same BFD session.
Total displayed BFD sessions	Total count of all the BFD sessions.

Sample Output

show spring-traffic-engineering lsp name

```
user@host> show spring-traffic-engineering lsp name lsp-name
```

To	State	LSP Name
10.1.1.7	Up	to-R1

show spring-traffic-engineering lsp detail

```
user@host> show spring-traffic-engineering lsp detail
```

```
10.1.1.7
  State: Up
  S-ERO: 24.1.1.1(80001) 10.1.1.3(4509) 11.2.1.2(9875)
  Bandwidth: 100M
  The above line is in IP address(label) format.
```

show spring-traffic-engineering overview

```
user@host> show spring-traffic-engineering overview
```

```
Overview of SPRING-TE:
  Route preference: 8
  Number of LSPs: 0 (Up: 0, Down: 0)
  External controllers:
    pccd
```

show spring-traffic-engineering sbfd detail

```
user@host> show spring-traffic-engineering sbfd detail
```

```
BFD name: V4-srte_bfd_session-1
BFD status: Down
Referencing LSPs:
  sr-lsp1:path1
  sr-lsp2:path1
SR-ERO hop count: 2
Hop 1 (Strict):
  NAI: IPv4 Adjacency ID, 1.2.1.1 -> 1.2.1.2
  SID type: 20-bit label, Value: 299776
Hop 2 (Strict):
```



```
NAI: IPv4 Adjacency ID, 2.3.0.1 -> 2.3.0.2
SID type: 20-bit label, Value: 299824
```

```
Total displayed BFD sessions: 2 (Up: 2, Down: 0)
```

show spring-traffic-engineering sbfd name <name> detail

```
user@host> show spring-traffic-engineering sbfd name V4-srte_bfd_session-1 detail
```

```
BFD name: V4-srte_bfd_session-1
BFD status: Up
Referencing LSPs:
  sr-lsp1:path1
  sr-lsp2:path1
SR-ERO hop count: 2
Hop 1 (Strict):
  NAI: IPv4 Adjacency ID, 1.2.1.1 -> 1.2.1.2
  SID type: 20-bit label, Value: 299776
Hop 2 (Strict):
  NAI: IPv4 Adjacency ID, 2.3.0.1 -> 2.3.0.2
  SID type: 20-bit label, Value: 299824
```