



Junos[®] OS for EX Series Ethernet Switches

MPLS for EX Series Switches

Release
12.3



Published: 2013-12-10

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

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

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

Junos[®] OS for EX Series Ethernet Switches MPLS for EX Series Switches
Release 12.3
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xiii
	Documentation and Release Notes	xiii
	Supported Platforms	xiii
	Using the Examples in This Manual	xiii
	Merging a Full Example	xiv
	Merging a Snippet	xiv
	Documentation Conventions	xv
	Documentation Feedback	xvi
	Requesting Technical Support	xvii
	Self-Help Online Tools and Resources	xvii
	Opening a Case with JTAC	xviii
Part 1	Overview	
Chapter 1	MPLS Overview	3
	Junos OS MPLS for EX Series Switches Overview	3
	Benefits of MPLS	4
	Additional Benefits of MPLS and Traffic Engineering	4
	Understanding Junos OS MPLS Components for EX Series Switches	5
	Provider Edge Switches	5
	MPLS Protocol and Label-Switched Paths	6
	Circuit Cross-Connect for Customer Edge Interfaces	6
	IP Over MPLS for Customer Edge Interfaces	7
	BGP for Layer 2 VPN and Layer 3 VPN Configurations (EX8200 and EX4500 Switches Only)	7
	Routing Instances for Layer 2 VPN and Layer 3 VPN (EX8200 and EX4500 Switches Only)	7
	Ethernet Encapsulation for Layer 2 VPN (EX8200 and EX4500 Switches Only)	7
	LDP for Layer 2 Circuits (EX8200 and EX4500 Switches Only)	7
	Provider Switch	7
	Components Required for All Switches in the MPLS Network	8
	Routing Protocol	8
	Traffic Engineering	8
	MPLS Protocol	8
	RSVP	9
	LDP	9
	Family mpls	9
	Planning Considerations While Using EX8200 Standalone Switches and EX8200 Virtual Chassis	10

	MPLS Support on EX4500 and EX4550 Standalone Switches and Virtual Chassis	10
	Understanding MPLS and Path Protection on EX Series Switches	11
	Understanding Using CoS with MPLS Networks on EX Series Switches	12
	EXP Classifiers and EXP rewrite Rules	12
	Guidelines for Using CoS Classifiers on CCCs	13
	Using CoS Classifiers with IP over MPLS	13
	Setting CoS Bits in an MPLS Header	14
	EXP Rewrite Rules	15
	Policer	15
	Schedulers	16
	Understanding MPLS Label Operations on EX Series Switches	17
	MPLS Label-Switched Paths and MPLS Labels on the Switches	17
	Reserved Labels	18
	MPLS Label Operations on the Switches	18
	Penultimate-Hop Popping and Ultimate-Hop Popping	19
	Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches	20
	MPLS-Based Layer 2 VPNs	20
	Layer 2 Circuits	21
	MPLS-Based Layer 3 VPNs	22
	Comparing an MPLS-Based Layer 3 VPN and an MPLS-Based Layer 2 VPN	22
	Understanding Using CoS with MPLS Networks on EX4500 and EX4550 Switches	23
	Ingress Provider Edge Switch	23
	Provider Switch	24
	Provider Switch That Is a Penultimate-Hop Pop Switch	24
	Egress Provider Edge Switch	24
Part 2	Configuration	
Chapter 2	Configuration Examples	29
	Example: Configuring MPLS on EX Series Switches	29
	Example: Combining CoS with MPLS on EX Series Switches	44
	Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches	56
	Example: Configuring MPLS-Based Layer 2 VPNs	65
Chapter 3	Configuration Tasks	79
	Configuring MPLS on Provider Switches (CLI Procedure)	80
	Configuring Path Protection in an MPLS Network (CLI Procedure)	81
	Configuring the Primary Path	83
	Configuring the Secondary Path	83
	Configuring the Revert Timer	84
	Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)	85
	Configuring the Ingress PE Switch	85
	Configuring the Egress PE Switch	87
	Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)	89

Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure)	92
Configuring CoS	92
Configuring an LSP Policer	93
Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure)	93
Configuring CoS	94
Configuring an LSP Policer	95
Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) . . .	96
Configuring CoS Bits for an MPLS Network (CLI Procedure)	97
Configuring an MPLS-Based Layer 2 VPN (CLI Procedure)	98
Configuring an MPLS-Based Layer 3 VPN (CLI Procedure)	101
Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN (CLI Procedure) . .	103
Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit (CLI Procedure)	106
Configuring an MPLS-Based VLAN CCC Using the Connection Method (CLI Procedure)	109
Configuring IPv6 Tunneling for MPLS (CLI Procedure)	110
Configuring Static Label Switched Paths for MPLS (CLI Procedure)	112
Configuring the Ingress PE Switch	112
Configuring the Provider and the Egress PE Switch	113
Configuring Bidirectional Forwarding Detection for MPLS (CLI Procedure)	114
Configuring BFD on Provider Edge and Provider Switches for an LDP-Based LSP	114
Configuring BFD on Provider Edge and Provider Switches for an RSVP-Based LSP	116
Chapter 4 Configuration Statements	119
[edit interfaces] Configuration Statement Hierarchy on EX Series Switches . . .	119
[edit protocols] Configuration Statement Hierarchy on EX Series Switches . . .	120
[edit protocols mpls] Configuration Statement Hierarchy on EX Series Switches	121
Supported Statements in the [edit protocols mpls] Hierarchy Level	121
Unsupported Statements in the [edit protocols mpls] Hierarchy Level	122
bfd-liveness-detection (Protocols MPLS)	133
connections (MPLS)	134
description (Protocols Layer 2 VPN)	134
encapsulation (Physical Interface)	135
encapsulation-type (Layer 2 VPNs)	140
exp	142
fec	143
instance-type	145
interface (MPLS)	147
l2circuit	148
l2vpn	149
label-switched-path	150
ldp	151
mpls	154
neighbor (Protocols Layer 2 Circuit)	156

path	157
periodic-traceroute	158
policing	160
primary	160
remote-interface-switch	161
remote-site-id	162
revert-timer	163
route-distinguisher	164
rsvp	166
secondary	167
signaling	168
site (Layer 2 Circuits)	169
site-identifier (Layer 2 Circuits)	170
standby	170
traffic-engineering	171
vrf-table-label	171
vrf-target	172

Part 3

Administration

Chapter 5

Routine Monitoring 175

Verifying That MPLS Is Working Correctly	175
Verifying the Physical Layer on the Switches	175
Verifying the Routing Protocol	176
Verifying the Core Interfaces Being Used for the MPLS Traffic	176
Verifying RSVP	176
Verifying the Assignment of Interfaces for MPLS Label Operations	177
Verifying the Status of the CCC	177
Verifying traceroute for Layer 3 VPN	178
Verifying Path Protection in an MPLS Network	178
Verifying the Primary Path	179
Verifying the RSVP-Enabled Interfaces	179
Verifying a Secondary Path	180

Chapter 6

Operational Commands 183

clear mpls lsp	184
clear rsvp session	186
clear rsvp statistics	188
ping mpls l2circuit	189
ping mpls l2vpn	192
ping mpls l3vpn	195
ping mpls ldp	198
ping mpls lsp-end-point	201
ping mpls rsvp	203
request mpls lsp adjust-autobandwidth	208
show connections	210
show link-management	213
show link-management peer	217
show link-management routing	219
show link-management statistics	222

show link-management te-link	224
show mpls admin-groups	226
show mpls call-admission-control	227
show mpls cspf	229
show mpls diffserv-te	231
show route forwarding-table	233
show mpls interface	240
show mpls interface	241
show mpls lsp	243
show mpls path	256
show rsvp interface	257
show rsvp neighbor	262
show rsvp session	267
show rsvp session	272
show rsvp statistics	281
show rsvp version	285
show ted database	288
show ted link	292
show ted protocol	294

List of Figures

Part 1	Overview	
Chapter 1	MPLS Overview	3
	Figure 1: Label Encoding	18
	Figure 2: MPLS Label Swapping	19
	Figure 3: Layer 2 VPN Connecting CE Switches	21
Part 2	Configuration	
Chapter 2	Configuration Examples	29
	Figure 4: Configuring MPLS on EX Series Switches	31
	Figure 5: MPLS-Based Layer 3 VPN	57
	Figure 6: MPLS-Based Layer 2 VPN	67

List of Tables

	About the Documentation	xiii
	Table 1: Notice Icons	xv
	Table 2: Text and Syntax Conventions	xv
Part 1	Overview	
Chapter 1	MPLS Overview	3
	Table 3: MPLS CoS Values	14
Part 2	Configuration	
Chapter 2	Configuration Examples	29
	Table 4: Components of the Ingress PE Switch in the Topology for MPLS with Interface-Based CCC	31
	Table 5: Components of the Egress PE Switch in the Topology for MPLS with Interface-Based CCC	32
	Table 6: Components of the Provider Switch in the Topology for MPLS with Interface-Based CCC	33
	Table 7: CoS Configuration Components on the Ingress PE Switch	45
	Table 8: CoS Configuration Components of the Egress PE Switch	46
	Table 9: CoS Configuration Components of the Provider Switch	46
	Table 10: Local CE Switch in the MPLS-Based Layer 3 VPN Topology	57
	Table 11: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology	57
	Table 12: Layer 3 VPN Components of the Local PE Switch	58
	Table 13: Layer 3 VPN Components of the Remote PE Switch	59
	Table 14: Local CE Routing Device in the MPLS-Based Layer 2 VPN Topology	67
	Table 15: Remote CE Routing Device in the MPLS-Based Layer 2 VPN Topology	67
	Table 16: Layer 2 VPN Components of the Local PE Routing Device	68
	Table 17: Layer 2 VPN Components of the Remote PE Routing Device	68
Chapter 4	Configuration Statements	119
	Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches	123
Part 3	Administration	
Chapter 6	Operational Commands	183
	Table 19: show connections Output Fields	211
	Table 20: show link-management Output Fields	213
	Table 21: show link-management peer Output Fields	217
	Table 22: show link-management routing Output Fields	219

Table 23: show link-management statistics Output Fields	222
Table 24: show link-management te-link Output Fields	224
Table 25: show mpls admin-groups Output Fields	226
Table 26: show mpls call-admission-control Output Fields	227
Table 27: show mpls cspf Output Fields	229
Table 28: show mpls diffserv-te Output Fields	231
Table 29: show route forwarding-table Output Fields	234
Table 30: show mpls interface Output Fields	240
Table 31: show mpls interface Output Fields	241
Table 32: show mpls lsp Output Fields	245
Table 33: show mpls path Output Fields	256
Table 34: show rsvp interface Output Fields	257
Table 35: show rsvp neighbor Output Fields	262
Table 36: show rsvp session Output Fields	268
Table 37: show rsvp session Output Fields	273
Table 38: show rsvp statistics Output Fields	281
Table 39: show rsvp version Output Fields	285
Table 40: show ted database Output Fields	288
Table 41: show ted link Output Fields	292
Table 42: show ted protocol Output Fields	294

About the Documentation

- Documentation and Release Notes on page xiii
- Supported Platforms on page xiii
- Using the Examples in This Manual on page xiii
- Documentation Conventions on page xv
- Documentation Feedback on page xvi
- Requesting Technical Support on page xvii

Documentation and Release Notes

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

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

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

Supported Platforms

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

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

Table 2 on page xv 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		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [MPLS Overview on page 3](#)

CHAPTER 1

MPLS Overview

- [Junos OS MPLS for EX Series Switches Overview on page 3](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)
- [Understanding MPLS and Path Protection on EX Series Switches on page 11](#)
- [Understanding Using CoS with MPLS Networks on EX Series Switches on page 12](#)
- [Understanding MPLS Label Operations on EX Series Switches on page 17](#)
- [Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches on page 20](#)
- [Understanding Using CoS with MPLS Networks on EX4500 and EX4550 Switches on page 23](#)

Junos OS MPLS for EX Series Switches Overview

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. See *EX Series Switch Software Features Overview* for a complete list of the Junos OS MPLS features that are supported on specific switches.



NOTE: MPLS configurations on the switches do not support:

- Q-in-Q tunneling

This topic describes:

- [Benefits of MPLS on page 4](#)
- [Additional Benefits of MPLS and Traffic Engineering on page 4](#)

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

- [Understanding MPLS Label Operations on EX Series Switches on page 17](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)
- [Understanding Using CoS with MPLS Networks on EX Series Switches on page 12](#)
- [Example: Configuring MPLS on EX Series Switches on page 29](#)
- [Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\) on page 103](#)
- [Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\) on page 106](#)

Understanding Junos OS MPLS Components for EX Series Switches

Juniper Networks Junos operating system (Junos OS) MPLS for Juniper Networks EX Series Ethernet Switches includes a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application or the specific switch. 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.

This topic includes:

- [Provider Edge Switches on page 5](#)
- [Provider Switch on page 7](#)
- [Components Required for All Switches in the MPLS Network on page 8](#)
- [Planning Considerations While Using EX8200 Standalone Switches and EX8200 Virtual Chassis on page 10](#)
- [MPLS Support on EX4500 and EX4550 Standalone Switches and Virtual Chassis on page 10](#)

Provider Edge Switches

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

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

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

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

- [MPLS Protocol and Label-Switched Paths on page 6](#)
- [Circuit Cross-Connect for Customer Edge Interfaces on page 6](#)
- [IP Over MPLS for Customer Edge Interfaces on page 7](#)
- [BGP for Layer 2 VPN and Layer 3 VPN Configurations \(EX8200 and EX4500 Switches Only\) on page 7](#)
- [Routing Instances for Layer 2 VPN and Layer 3 VPN \(EX8200 and EX4500 Switches Only\) on page 7](#)

- [Ethernet Encapsulation for Layer 2 VPN \(EX8200 and EX4500 Switches Only\)](#) on page 7
- [LDP for Layer 2 Circuits \(EX8200 and EX4500 Switches Only\)](#) on page 7

MPLS Protocol and Label-Switched Paths

Each PE switch must be configured to support the MPLS protocol. The configuration of a label-switched path (LSP) depends upon which signaling protocol is used:

- If the RSVP signaling protocol is used, the LSPs must be explicitly configured at the **[edit protocols mpls]** hierarchy level.
- If the LDP signaling protocol is used, LSP configuration is not required. (LDP signaling is used with Layer 2 circuit configurations.)

Circuit Cross-Connect for Customer Edge Interfaces

You can configure the customer edge interface of the PE switches as a circuit cross-connect (CCC) to create a transparent connection between two circuits. When you configure an interface as a CCC, the interface is removed from the default VLAN if it was a member of that VLAN. The interface becomes an MPLS tunnel—used exclusively for MPLS packets. You can create different CCCs for different customers or for segregating different traffic streams over different MPLS tunnels.

Using a CCC configuration, you can connect the following types of interfaces:

- A local interface with a remote interface or VLAN
- A local VLAN with a remote interface or VLAN



NOTE: To configure a VLAN circuit as a CCC, you must enable VLAN tagging and specify a VLAN ID. You must specify the same VLAN ID on both ends of the CCC.

The VLAN CCC configuration must use the same type of switch for both PE switches. For example, you cannot use an EX8200 switch for one PE switch and an EX3200, EX4200, or EX4500 switch for the other PE switch.

MPLS on EX Series switches does not support the following types of CCC configurations:

- Routed VLAN interfaces (RVIs) on switches other than EX8200 switches
- Q-in-Q tunneling

On EX8200 switches only, the following types of CCC configurations are supported:

- Local switching—Connecting interfaces on the same switch
- MPLS tunneling—Using LSPs as the conduit to connect two distant interface circuits

- LSP stitching—Connecting LSPs that fall into two different traffic engineering database areas
- RVIs on customer edge interfaces

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 Using IP Over MPLS (CLI Procedure)” on page 85.

BGP for Layer 2 VPN and Layer 3 VPN Configurations (EX8200 and EX4500 Switches Only)

If you are implementing a Layer 2 virtual private network (VPN) or a Layer 3 VPN, you must configure the BGP routing protocol on the PE switches.

Routing Instances for Layer 2 VPN and Layer 3 VPN (EX8200 and EX4500 Switches Only)

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.

EX Series switches support the following types of routing instances:

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

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

Ethernet Encapsulation for Layer 2 VPN (EX8200 and EX4500 Switches Only)

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.

LDP for Layer 2 Circuits (EX8200 and EX4500 Switches Only)

If you are implementing a Layer 2 circuit configuration, you must configure LDP as the signaling protocol on the PE switches.

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

Components Required for All Switches in the MPLS Network

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

- [Routing Protocol on page 8](#)
- [Traffic Engineering on page 8](#)
- [MPLS Protocol on page 8](#)
- [RSVP on page 9](#)
- [LDP on page 9](#)
- [Family mpls on page 9](#)

Routing Protocol

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

The core 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 core 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. The configuration of traffic engineering depends upon which routing protocol is being used:

- With OSPF—Traffic engineering needs to be enabled.
- With IS-IS—Traffic engineering is enabled by default.

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 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 switch receives the path message, it sends a reservation message back to the ingress switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress 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 the Junos OS and is not in the packet-forwarding path.

LDP

LDP is a signaling protocol available on EX8200 switches and EX4500 switches. LDP is a simple, fast-acting signaling protocol that automatically establishes LSP adjacencies within an MPLS network. Switches 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 true traffic engineering like RSVP. LDP does not support bandwidth reservation or traffic constraints.



NOTE: LDP can be used with basic MPLS or with MPLS and a Layer 2 circuit configuration.

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.

Planning Considerations While Using EX8200 Standalone Switches and EX8200 Virtual Chassis

EX8200 standalone switches and EX8200 Virtual Chassis support up to 4000 MPLS tunnel start entries. MPLS tunnel start entries define tunnels to be used to transmit packets across the MPLS network. MPLS tunnel start entries are required for label manipulation (push and swap) beyond two labels. In scenarios requiring make-before-break, for example Layer 3 VPN packets on an ingress PE switch transmitted from one LSP to another, the system creates new tunnel start entries even before the old tunnel start entries are deleted. In such cases, consider a maximum of 2000 MPLS tunnel start entries while planning your MPLS implementation.

MPLS Support on EX4500 and EX4550 Standalone Switches and Virtual Chassis

EX4500 standalone switches and EX4500 Virtual Chassis now support all MPLS features that are supported on EX8200 switches with the following exceptions:

- MPLS is not supported in a mixed EX4200 and EX4500 Virtual Chassis.
- IP over MPLS is not supported when an EX4500 switch is positioned as a non-penultimate hop popping (PHP) MPLS switch; that is, when the label operation is swap. However, EX4500 switches support CCC, Layer2 VPNs, Layer2 circuits, and Layer3 VPNs the same way these are supported on EX8200 switches.
- MPLS over RVIs, LSP statistics, unicast reverse-path forwarding (RPF) statistics, MPLS class of service (CoS), traffic policing, Diffserv-aware LSPs, graceful Routing Engine switchover (GRES), and equal-cost multipath (ECMP) are not supported.
- LSP ping and traceroute for CCCs, Layer 2 circuits, and Layer 2 VPNs are not supported.
- An MPLS configuration that consists of a mix of EX8200 and EX4500 switches does not support VLAN-CCCs.
- VLAN-CCCs require that the VLAN ID is the same at both ends of the connection. The VLAN ID translation feature is not supported.
- EX4500 standalone switches and EX4500 Virtual Chassis support a maximum of 125 instances of Layer 2 VPN, Layer 3 VPN, or CCC connections; or a combination of these.



NOTE: EX4500 switches do not support non-stop routing (NSR) for MPLS. With NSR support for MPLS, EX4500 switches can support up to 225 simultaneous instances of Layer 2 VPN or Layer 3 VPN or CCC connections.

- Time to live (TTL) of MPLS packets is not decremented in the ingress MPLS switch.
- The pipe model of TTL handling is not supported on a Layer 3 VPN if an EX4500 switch is configured as the ingress provider edge (PE) switch.

- Related Documentation**
- [Understanding MPLS and Path Protection on EX Series Switches on page 11](#)
 - [Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches on page 20](#)
 - [Example: Configuring MPLS on EX Series Switches on page 29](#)
 - [Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
 - [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
 - [Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\) on page 103](#)
 - [Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\) on page 106](#)
 - [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)
 - [Junos OS VPNs Configuration Guide](#)
 - [Junos OS MPLS Applications Configuration Guide](#)

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 81.

- Related Documentation**
- [Junos OS MPLS for EX Series Switches Overview on page 3](#)
 - [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)
 - [Example: Configuring MPLS on EX Series Switches on page 29](#)

- [Configuring MPLS on Provider Edge Switches \(CLI Procedure\)](#)
- [Junos OS MPLS Applications Configuration Guide](#)

Understanding Using CoS with MPLS Networks on EX Series Switches

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 on page 12](#)
- [Guidelines for Using CoS Classifiers on CCCs on page 13](#)
- [Using CoS Classifiers with IP over MPLS on page 13](#)
- [Setting CoS Bits in an MPLS Header on page 14](#)
- [EXP Rewrite Rules on page 15](#)
- [Policer on page 15](#)
- [Schedulers on page 16](#)

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 *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.

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 *Junos OS Class of Service Configuration Guide*.



NOTE: Configuring the PLP drop profile to drop packets more aggressively (for example, setting the CoS value from 6 to 7) decreases the likelihood of traffic getting through.

Table 3 on page 14 summarizes how MPLS CoS values correspond to the transmit queue and PLP bit. Note that in MPLS, the mapping between the CoS bit value and the output queue is hard-coded. You cannot configure the mapping for MPLS; you can configure it only for IPv4 traffic flows, as described in the *Junos OS Class of Service Configuration Guide*.

Table 3: 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

Table 3: MPLS CoS Values (*continued*)

MPLS CoS Value	Bits	Transmit Queue	PLP Bit
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 \(CLI Procedure\) on page 93](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) on page 92](#)
- *Configuring Rewrite Rules for EXP Classifiers on MPLS Networks (CLI Procedure)*
- [Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) on page 96](#)
- [Configuring CoS Bits for an MPLS Network \(CLI Procedure\) on page 97](#)

Understanding MPLS Label Operations on EX Series Switches

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 on page 17](#)
- [Reserved Labels on page 18](#)
- [MPLS Label Operations on the Switches on page 18](#)
- [Penultimate-Hop Popping and Ultimate-Hop Popping on page 19](#)

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 1 on page 18](#) shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 1: Label Encoding



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 2 on page 19 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 2: MPLS Label Swapping




Figure 2 on page 19 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

- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)
- [Example: Configuring MPLS on EX Series Switches on page 29](#)
- [Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
- [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
- [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)
- [Junos OS VPNs Configuration Guide](#)
- [Junos OS MPLS Applications Configuration Guide](#)

Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches

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 on page 20](#)
- [Layer 2 Circuits on page 21](#)
- [MPLS-Based Layer 3 VPNs on page 22](#)
- [Comparing an MPLS-Based Layer 3 VPN and an MPLS-Based Layer 2 VPN on page 22](#)

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.

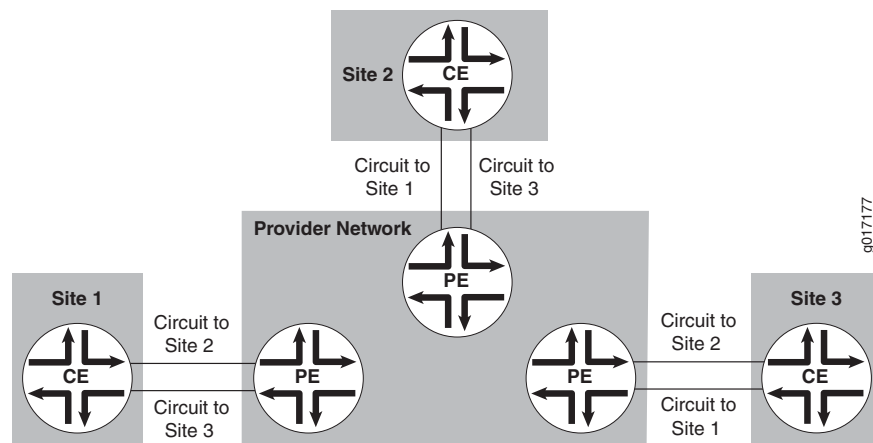
On a Layer 2 VPN, routing occurs on the customer's switches, typically on the CE switch. The CE switch connected to a service provider on a Layer 2 VPN must select the

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 3 on page 21](#) 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 3: 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 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.

Comparing an MPLS-Based Layer 3 VPN and an MPLS-Based Layer 2 VPN

EX8200 and EX4500 switches can support the following kinds of MPLS-based VPNs:

- Layer 3 VPNs—The service provider participates in the customer's Layer 3 routing. Layer 3 VPNs allow customers to leverage the service provider's technical expertise to ensure efficient site-to-site routing. The customer's CE switch uses a routing protocol such as BGP or OSPF to communicate with the provider's PE switch to carry IP prefixes across the network. MPLS-based Layer 3 VPNs use IP over MPLS. Other protocol packets are not supported.
- Layer 2 VPNs—The service provider interconnects customer sites using Layer 2 technology. Layer 2 VPNs give customers complete control over their own routing.

Related Documentation

- [Understanding MPLS Label Operations on EX Series Switches on page 17](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)
- [Example: Configuring MPLS-Based Layer 2 VPNs on page 65](#)
- [Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches on page 56](#)
- [Junos OS VPNs Configuration Guide](#)

- *Junos OS MPLS Applications Configuration Guide*

Understanding Using CoS with MPLS Networks on EX4500 and EX4550 Switches

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 Juniper Networks Junos operating system (Junos OS) MPLS features that are supported on specific Juniper Networks EX Series Ethernet Switches..

Juniper Networks EX4500 and EX4550 Ethernet Switches support DiffServ code point (DSCP) or IP precedence and IEEE 802.1p CoS classifiers on the customer-edge (CE) interfaces of the ingress provider edge (PE) switch. DSCP or IP precedence classifiers are used for Layer 3 packets. IEEE 802.1p classifiers are used for Layer 2 packets. EX4500 and EX4550 switches also support DSCP or IP precedence and IEEE 802.1p CoS rewrite rules on the CE interfaces of the egress PE switch.

EX4500 and EX4550 switches enable a default EXP classifier and a default EXP rewrite rule. If you want, you can also configure a custom EXP classifier and a custom EXP rewrite rule. However, the switches support only one type of EXP classifier—default or custom—and only one type of EXP rewrite rule—default or custom. The default and custom EXP classifier or rewrite rules are global to the switch or a virtual chassis and not applied on any interface.

After CoS is configured, each incoming IP packet is processed as follows at each switch in your MPLS network:

- [Ingress Provider Edge Switch on page 23](#)
- [Provider Switch on page 24](#)
- [Provider Switch That Is a Penultimate-Hop Pop Switch on page 24](#)
- [Egress Provider Edge Switch on page 24](#)

Ingress Provider Edge Switch

On the ingress PE switch:

- Configure and bind a classifier to the CE interface. The choice of classifier to use—DSCP, IP, or IEEE 802.1p—depends on the type of IP packet. Using these classifiers, the switch assigns a forwarding class and a loss priority value to the incoming IP packet on the basis of the packet's DSCP, IP precedence, or IEEE 802.1p value. This configuration applies to both the inet and the circuit cross-connect (CCC families and also to VLAN CCC interfaces)..
- EXP EXP re-marking in MPLS packets on the core facing interfaces is achieved through the Global CoS profile table, which can be configured by the user. You cannot control EXP rewrite on individual interfaces. The EXP value is set on the basis of the forwarding class and loss priority that were assigned to the packet. The rewrite applies to both outer and inner labels. After a rewrite rule is configured, it is applied to all the MPLS packets entering the tunnel.

- The IP precedence bits of DSCP are not copied to the EXP bits. EXP bits are always taken from the Global CoS profile table.

For more details about configure CoS on an MPLS PE switch, see [“Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\)” on page 93](#) or [“Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\)” on page 92](#).

Provider Switch

On the provider switch:

- The EXP classification occurs through the Global EXP to CoS profile table, which can be configured by the user. You cannot control EXP classification on individual interfaces. Using this table, a forwarding class and a loss priority value is assigned to the incoming MPLS packet on the basis of the EXP value of the outer MPLS label.
- The EXP value of the outer label is not preserved during a SWAP operation at the provider switch. When the outermost label is swapped with a new label, the EXP value in that label is also reconfigured on the basis of the Global CoS profile table.

To retain the EXP value from the outer label, you must configure an EXP classifier to map the label's EXP to a CoS profile and also configure an EXP rewrite rule that maps the CoS profile to the same EXP value.

- EXP rewrite is achieved through the Global CoS profile table, which can be configured by the user. You cannot control EXP rewrite on individual interfaces. The EXP value is set on the basis of the forwarding class and loss priority that were assigned to the packet. The rewrite applies to only the outer label. After a rewrite rule is configured, it is applied to all the MPLS packets entering the tunnel.

For more details on configuring CoS on an MPLS provider switch, see [“Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\)” on page 96](#).

Provider Switch That Is a Penultimate-Hop Pop Switch

On the provider switch, that is a Penultimate-hop Pop switch:

- The EXP classification occurs through the Global EXP to CoS profile table, which can be configured by the user. You cannot control EXP classification on individual interfaces. Using this table, a forwarding class and a loss priority value is assigned to the incoming MPLS packet on the basis of the EXP value of the outer MPLS label.
- On popping the outer label, the EXP value is not copied from the outer label to the inner label and you cannot use the Junos OS CLI to rewrite the inner label's EXP value.

Egress Provider Edge Switch

On the egress provider edge switch:

- The EXP classification occurs through the Global EXP to CoS profile table, which can be configured by the user. You cannot control EXP classification on individual interfaces. Using this table, a forwarding class and a loss priority value is assigned to the incoming MPLS packet on the basis of the EXP value of the MPLS label.

- The EXP value is not copied to the IP precedence bits, but it is possible to overwrite the DSCP, IP precedence, or IEEE 802.1p values by configuring and binding a rewrite rule to the CE interface for Layer 3 VPN. Using these rules, the DSCP, IP precedence, or IEEE 802.1p value is overwritten on the basis of the forwarding class and loss priority that were assigned to the packet.
- DSCP, IP precedence, and IEEE 802.1p rewrite rules are not supported on CCC or VLAN CCC interfaces.

For more details about configure CoS on an MPLS PE switch, see [“Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\)” on page 93](#) or [“Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\)” on page 92](#).

**Related
Documentation**

- *Understanding CoS Classifiers*
- *Example: Configuring CoS on EX Series Switches*
- [Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) on page 93](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) on page 92](#)
- [Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) on page 96](#)
[Configuring CoS Bits for an MPLS Network \(CLI Procedure\) on page 97](#)

PART 2

Configuration

- [Configuration Examples on page 29](#)
- [Configuration Tasks on page 79](#)
- [Configuration Statements on page 119](#)

CHAPTER 2

Configuration Examples

- [Example: Configuring MPLS on EX Series Switches on page 29](#)
- [Example: Combining CoS with MPLS on EX Series Switches on page 44](#)
- [Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches on page 56](#)
- [Example: Configuring MPLS-Based Layer 2 VPNs on page 65](#)

Example: Configuring MPLS on EX Series Switches

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 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 103](#) or [“Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\)” on page 106](#).

- [Requirements on page 30](#)
- [Overview and Topology on page 30](#)
- [Configuring the Local PE Switch on page 34](#)
- [Configuring the Remote PE Switch on page 37](#)
- [Configuring the Provider Switch on page 39](#)
- [Verification on page 41](#)

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 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 4 on page 31](#) shows the topology used in this example.

Figure 4: Configuring MPLS on EX Series Switches




Table 4 on page 31 shows the MPLS configuration components used for the ingress PE switch in this example.

Table 4: 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.
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 .

Table 4: Components of the Ingress PE Switch in the Topology for MPLS with Interface-Based CCC (continued)

Property	Settings	Description
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 5 on page 32](#) shows the MPLS configuration components used for the egress PE switch in this example.

Table 5: 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.
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.

Table 5: Components of the Egress 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 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 6 on page 33](#) shows the MPLS configuration components used for the provider switch in this example.

Table 6: 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.

Table 6: Components of the Provider Switch in the Topology for MPLS with Interface-Based CCC (*continued*)

Property	Settings	Description
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 127.1.1.32
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 127.1.1/32
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 103 or “Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit (CLI Procedure)” on page 106.

```
[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

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

- [Verifying the Physical Layer on the Switches on page 41](#)
- [Verifying the Routing Protocol on page 42](#)
- [Verifying the Core Interfaces Being Used for MPLS Traffic on page 42](#)
- [Verifying the Status of the RSVP Sessions on page 42](#)
- [Verifying the Assignment of Interfaces for MPLS Label Operations on page 43](#)
- [Verifying the Status of the CCC on page 43](#)

### 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  | Legend for circuit types            |
| Up -- operational              | intf -- interface                   |
| RmtDn -- remote CCC down       | tlsp -- transmit LSP                |
| Restart -- restarting          | rlsp -- receive LSP                 |

| Connection/Circuit | Type   | St | Time last up    | # Up trans |
|--------------------|--------|----|-----------------|------------|
| ge1-to-pe2         | rmt-if | Up | Feb 17 05:00:09 | 1          |
| ge-0/0/1.0         | intf   | Up |                 |            |
| lsp_to_pe1_ge1     | tlsp   | Up |                 |            |
| lsp_to_pe2_ge1     | 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 Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
- [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
- [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)
- [Junos OS MPLS for EX Series Switches Overview on page 3](#)

---

## Example: Combining CoS with MPLS on EX Series Switches

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 on page 44](#)
- [Overview and Topology on page 45](#)
- [Configuring the Local PE Switch on page 47](#)
- [Configuring the Remote PE Switch on page 49](#)
- [Configuring the Provider Switch on page 49](#)
- [Verification on page 50](#)

## 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 EX Series Switches” on page 29](#). 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 EX Series Switches](#)” on page 29) 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 7 on page 45](#) shows the CoS configuration components added to the ingress PE switch.

**Table 7: 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. | <b>policing filter mypolicer</b> | Name of the rate-limiting policer.               |
|                                                    | <b>filter myfilter</b>           | Name of the filter, which refers to the policer  |
| Custom DSCP classifier                             | <b>dscp1</b>                     | Specifies the name of the custom DSCP classifier |
| Custom EXP rewrite rule                            | <b>e1</b>                        | Name of the custom EXP rewrite rule.             |

**Table 7: CoS Configuration Components on the Ingress PE Switch (*continued*)**

| Property                | Settings                         | Description                                                                                                                                             |
|-------------------------|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| Customer-edge interface | <b>ge-0/0/1.0</b>                | Interface that receives packets from devices outside the network.<br><br>The custom DSCP classifier must be specified on this CCC interface.            |
| Core interfaces         | <b>ge-0/0/5.0 and ge-0/0/6.0</b> | Interfaces that transmit MPLS packets to other switches within the MPLS network.<br><br>The EXP rewrite rule is applied implicitly to these interfaces. |

[Table 8 on page 46](#) shows the CoS configuration components added to the egress PE switch in this example.

**Table 8: CoS Configuration Components of the Egress PE Switch**

| Property                             | Settings                         | Description                                                                                                                                                                    |
|--------------------------------------|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote provider edge switch hardware | EX Series switch                 | PE-2                                                                                                                                                                           |
| Custom EXP classifier                | <b>exp1</b>                      | Name of custom EXP classifier                                                                                                                                                  |
| Customer-edge interface              | <b>ge-0/0/1.0</b>                | 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                      | <b>ge-0/0/7.0 and ge-0/0/8.0</b> | 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 9 on page 46](#) shows the MPLS configuration components used for the provider switch in this example.

**Table 9: 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    | <b>exp1</b>      | Name of the custom EXP classifier.                    |
| Custom EXP rewrite rule  | <b>e1</b>        | Name of the custom EXP rewrite rule.                  |



Table 9: CoS Configuration Components of the Provider Switch (*continued*)

| Property                                                                        | Settings                                | Description                                                                                                                                                                                         |
|---------------------------------------------------------------------------------|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Core interfaces receiving packets from other MPLS switches.                     | <b>ge-0/0/5.0</b> and <b>ge-0/0/6.0</b> | 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.                     |
| Core interfaces transmitting packets to other switches within the MPLS network. | <b>ge-0/0/7.0</b> and <b>ge-0/0/8.0</b> | 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

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

- [Verifying That the Policer Firewall Filter Is Operational on page 51](#)
- [Verifying That the CoS Classifiers Are Going to the Right Queue on page 51](#)
- [Verifying the CoS Forwarding Table Mapping on page 54](#)
- [Verifying the Rewrite Rules on page 55](#)

### 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/<br>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 Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
  - [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
  - [Understanding Using CoS with MPLS Networks on EX Series Switches on page 12](#)
  - [Monitoring CoS Forwarding Classes](#)

## Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches

---

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 on page 56](#)
- [Overview and Topology on page 56](#)
- [Configuring the Local PE Switch on page 59](#)
- [Configuring the Remote PE Switch on page 62](#)
- [Verification on page 64](#)

### 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 85.
- Configure one or more provider switches. See “[Configuring MPLS on Provider Switches \(CLI Procedure\)](#)” on page 80.



**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 5 on page 57 illustrates the topology of this MPLS-based Layer 3 VPN.

**Figure 5: MPLS-Based Layer 3 VPN**




Table 10 on page 57 shows the settings of the customer edge interface on the local CE switch.

**Table 10: Local CE Switch in the MPLS-Based Layer 3 VPN Topology**

| Property                 | Settings                                                                      | Description                         |
|--------------------------|-------------------------------------------------------------------------------|-------------------------------------|
| Local CE switch hardware | EX8200 switch                                                                 | CE1                                 |
| Customer edge interface  | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.14/16</b> | Interface that connects CE1 to PE1. |

Table 11 on page 57 shows the settings of the customer edge interface on the remote CE switch.

**Table 11: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology**

| Property                  | Settings      | Description |
|---------------------------|---------------|-------------|
| Remote CE switch hardware | EX8200 switch | CE2         |

Table 11: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology (*continued*)

| Property                | Settings                                                                      | Description                         |
|-------------------------|-------------------------------------------------------------------------------|-------------------------------------|
| Customer edge interface | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 11.22.26.1/16</b> | Interface that connects CE2 to PE2. |

Table 12 on page 58 shows the Layer 3 VPN components of the local PE switch.

Table 12: Layer 3 VPN Components of the Local PE Switch

| Property                 | Settings                                                                                                                        | Description                                                                                                                                                                                                                                                                       |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local PE switch hardware | EX8200 switch                                                                                                                   | PE1                                                                                                                                                                                                                                                                               |
| Customer edge interface  | <b>ge-5/0/24 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.1/16</b>                                                    | Connects PE1 to CE1.<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface           | <b>xe-6/0/0 unit 0</b><br><b>family inet address 60.0.0.60/16</b><br><b>family iso;</b><br><b>family mpls</b>                   | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                        |
| Loopback interface       | <b>lo0 unit 0</b><br><b>family inet address 21.21.21.21/32</b><br><b>family iso address</b><br><b>49.0001.2102.1021.0210.00</b> | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                  |
| BGP                      | <b>bgp</b>                                                                                                                      | Added for the Layer 3 VPN configuration.                                                                                                                                                                                                                                          |
| Routing instance         | <b>L3VPN-1</b>                                                                                                                  | Added for the Layer 3 VPN configuration.                                                                                                                                                                                                                                          |

Table 13 on page 59 shows the Layer 3 VPN components of the remote PE switch.

**Table 13: Layer 3 VPN Components of the Remote PE Switch**

| Property                  | Settings                                                                                                                        | Description                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote PE switch hardware | EX8200 switch                                                                                                                   | PE2                                                                                                                                                                                                                                                                                                                                                    |
| Customer edge interface   | <b>ge-11/0/14 unit 0</b><br><b>family inet</b><br><b>address 11.22.26.14/16</b><br><b>family mpls</b>                           | Connects PE2 to CE2.<br><br>For the Layer 3 VPN configuration, added <b>family mpls</b> .<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface            | <b>xe-6/0/0/ unit 0</b><br><b>family inet address 60.2.0.60/16</b><br><b>family iso</b><br><b>family mpls</b>                   | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                                                             |
| Loopback interface        | <b>lo0 unit 0</b><br><b>family inet address 22.22.22.22/32</b><br><b>family iso address</b><br><b>49.0001.2202.1022.0220.00</b> | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                                                                                       |
| BGP                       | <b>bgp</b>                                                                                                                      | Added for the Layer 3 VPN configuration.                                                                                                                                                                                                                                                                                                               |
| Routing instances         | <b>L3VPN-1</b>                                                                                                                  | 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

To confirm that the MPLS-based Layer 3 VPN is working properly, perform these tasks:

- [Verifying Peering and Adjacency on page 64](#)
- [Verifying That the Local CE Switch Can Ping the Local PE Switch on page 65](#)
- [Verifying That the Local PE Switch Can Ping the Local CE Switch on page 65](#)

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 \(CLI Procedure\) on page 85](#)
- [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)

Example: Configuring MPLS-Based Layer 2 VPNs

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 on page 66](#)
- [Overview and Topology on page 66](#)
- [Configuring the Local PE Routing Device on page 69](#)
- [Configuring the Remote PE Routing Device on page 72](#)
- [Verification on page 74](#)

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 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 89](#).
- Configure one or more provider devices. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)” on page 80](#).



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 6 on page 67 illustrates the topology of this MPLS-based Layer 2 VPN.

Figure 6: MPLS-Based Layer 2 VPN




Table 14 on page 67 shows the settings of the customer edge interface on the local CE device.

Table 14: 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 11.0.0.2/16	Interface that connects CE1 to PE1.

Table 15 on page 67 shows the settings of the customer edge interface on the remote CE routing device.

Table 15: Remote CE Routing Device in the MPLS-Based Layer 2 VPN Topology

Property	Settings	Description
Remote CE routing device hardware	Routing device	CE2

Table 15: Remote CE Routing Device in the MPLS-Based Layer 2 VPN Topology (*continued*)

Property	Settings	Description
Customer edge interface	ge-0/0/0 unit 0 family inet address 11.0.0.1/16	Interface that connects CE2 to PE2.

Table 16 on page 68 shows the Layer 2 VPN components of the local PE routing device.

Table 16: 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 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 49.0001.2102.2021.0210.00	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.
BGP	bgp	Added for the Layer 2 VPN configuration.
Routing instance	vpn1	Added for the Layer 2 VPN configuration

Table 17 on page 68 shows the Layer 2 VPN components of the remote PE routing device.

Table 17: Layer 2 VPN Components of the Remote PE Routing Device

Property	Settings	Description
PE routing device hardware	Routing device	PE2

Table 17: Layer 2 VPN Components of the Remote PE Routing Device (*continued*)

Property	Settings	Description
Customer edge interface	<code>ge-11/0/0</code> <code>encapsulation ethernet-ccc</code> <code>unit 0</code> <code>family ccc</code>	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	<code>xe-6/0/0</code> <code>unit 0</code> <code>family inet</code> <code>address 60.2.0.61/16</code> <code>family iso</code> <code>family mpls</code>	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	<code>lo0</code> <code>unit 0</code> <code>family inet</code> <code>address 22.22.22.22/32</code> <code>family iso</code> <code>address 49.0001.2202.2022.0220.00</code>	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	<code>bgp</code>	Added for the Layer 2 VPN configuration.
Routing instance	<code>vpn1</code>	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 21.21.21.21 family l2vpn signaling
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 22.22.22.22
set routing-instances vpn1 instance-type l2vpn
set routing-instances vpn1 interface ge-5/0/0
set routing-instances vpn1 route-distinguisher 21.21.21.21
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 PE2"
set routing-instances vpn1 protocols l2vpn site JE-V21 site-identifier 21 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 21.21.21.21 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 22.22.22.22
```
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 21.21.21.21
```
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 PE2"
```
11. Configure the routing-instance protocols site:

```
[edit routing-instances]
user@PE1# set vpn1 protocols l2vpn site JE-V21 site-identifier 21remote-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 60.0.0.60/16;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 21.21.21.21/32;
      }
      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 22.22.22.22;
    }
    interface xe-0/0/6.0;
  }
  bgp {
    group ibgp
    type internal
    local-address 21.21.21.21
    family inet-vpn
    unicast
  }
}

routing-instances {
```

```

vpn1 {
  instance-type l2vpn;
  interface ge-5/0/0.0;
  route-distinguisher 21.21.21.21;
  vrf-target target:21:21;
  protocols {
    l2vpn {
      encapsulation-type ethernet;
      interface ge-5/0/0.0 {
        description "BETWEEN PE1 AND PE2";
      }
      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 22.22.22.22 family l2vpn signaling
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 21.21.21.21
set routing-instances vpn1 instance-type l2vpn
set routing-instances vpn1 interface ge-11/0/0
set routing-instances vpn1 route-distinguisher 21.21.21.21
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 PE2"
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 22.22.22.22 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 21.21.21.21
```
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 21.21.21.21:21
```
  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 PE2"
```
  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 60.2.0.61/16;
 }
 family mpls;
 }
 }
}
```

```
}
lo0 {
 unit 0 {
 family inet {
 address 22.22.22.22/32;
 }
 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 21.21.21.21;
 }
 }
 interface xe-0/0/6.0;
 bgp {
 group ibgp
 type internal
 local-address 21.21.21.21
 family inet-vpn
 unicast
 }
 routing-instances {
 vpn1 {
 instance-type l2vpn;
 interface ge-11/0/0.0;
 route-distinguisher 21.21.21.21:21;
 vrf-target target:21:21;
 protocols {
 l2vpn {
 encapsulation-type ethernet;
 interface ge-11/0/0.0 {
 description "BETWEEN PE1 AND PE2";
 }
 site T26-VPN1 {
 site-identifier 26;
 interface ge-11/0/0.0 {
 remote-site-id 21;
 }
 }
 }
 }
 }
 }
}
```

## Verification

To confirm that the MPLS-based Layer 2 VPN is working properly, perform these tasks:

- [Verifying the Layer 2 VPN Connection on page 75](#)
- [Verifying the Status of MPLS Label-Switched Paths on page 75](#)

- [Verifying BGP Status on page 76](#)
- [Verifying the Status of the RSVP Sessions on page 76](#)
- [Verifying the Routes in the Routing Table on page 77](#)
- [Pinging the Layer 2 VPN Connections on page 78](#)

### 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: 22.22.22.22, 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 (22.22.22.22)** 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    |
|-------------|-------------|-------|----|---|------------|------------|
| 22.22.22.22 | 21.21.21.21 | Up    | 0  | * |            | lsp_to_pe2 |

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

| To          | From        | State | Rt | Style | Labelin | Labelout | LSPname    |
|-------------|-------------|-------|----|-------|---------|----------|------------|
| 21.21.21.21 | 22.22.22.22 | 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 (22.22.22.22)** is up, and the Egress LSP from the remote PE routing device to this PE routing device (**21.21.21.21**) 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.12vpn.0           | 1   | 1     | 0      | 0     | 0          | 0       | 0      | 0                                         | 0       |
| Peer                  | AS  | InPkt | OutPkt | OutQ  | Flaps      | Last    | Up/Dwn | State #Active/Received/Accepted/Damped... |         |
| 22.22.22.22           | 10  | 33    | 34     | 0     | 1          | 13:24   |        |                                           |         |
| Establ                |     |       |        |       |            |         |        |                                           |         |
| bgp.12vpn.0: 1/1/1/0  |     |       |        |       |            |         |        |                                           |         |
| vpn2.12vpn.0: 1/1/1/0 |     |       |        |       |            |         |        |                                           |         |

**Meaning** The output shows that the remote PE routing device (**22.22.22.22**) 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
22.22.22.22 21.21.21.21 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
21.21.21.21 22.22.22.22 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 PE1, 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
```

```
2:2:27:27/96
 *[BGP/170] 00:13:55, localpref 100, from 22.22.22.22
 AS path: I
 > to 60.2.0.24 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
```

```
2:2:27:27/96
 *[BGP/170] 00:14:00, localpref 100, from 22.22.22.22
 AS path: I
 > to 60.2.0.24 via ge-6/0/46.0, label-switched-path lsp_to_pe2
2:2:28:27/96
 *[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 EX Series Switches on page 29](#)
- [Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches on page 56](#)
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) on page 98](#)



## CHAPTER 3

# Configuration Tasks

- [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)
- [Configuring Path Protection in an MPLS Network \(CLI Procedure\) on page 81](#)
- [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
- [Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) on page 92](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect \(CLI Procedure\) on page 93](#)
- [Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) on page 96](#)
- [Configuring CoS Bits for an MPLS Network \(CLI Procedure\) on page 97](#)
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) on page 98](#)
- [Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) on page 101](#)
- [Configuring an MPLS-Based VLAN CCC Using a Layer 2 VPN \(CLI Procedure\) on page 103](#)
- [Configuring an MPLS-Based VLAN CCC Using a Layer 2 Circuit \(CLI Procedure\) on page 106](#)
- [Configuring an MPLS-Based VLAN CCC Using the Connection Method \(CLI Procedure\) on page 109](#)
- [Configuring IPv6 Tunneling for MPLS \(CLI Procedure\) on page 110](#)
- [Configuring Static Label Switched Paths for MPLS \(CLI Procedure\) on page 112](#)
- [Configuring Bidirectional Forwarding Detection for MPLS \(CLI Procedure\) on page 114](#)

## Configuring MPLS on Provider Switches (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 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 127.1.1/32
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 EX Series Switches on page 29](#)
- [Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
- [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) on page 98](#)
- [Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) on page 101](#)

## 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 85](#) or [“Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 89](#).
- Configured at least one provider (transit) switch. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)” on page 80](#).

- Verified the configuration of your MPLS network. See [“Verifying That MPLS Is Working Correctly”](#) on page 175.

To configure path protection, complete the following tasks on the ingress provider edge switch:

1. [Configuring the Primary Path on page 83](#)
2. [Configuring the Secondary Path on page 83](#)
3. [Configuring the Revert Timer on page 84](#)

## 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 on page 11](#)

## 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 Provider Switches \(CLI Procedure\)” on page 80](#).

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. [Configuring the Ingress PE Switch on page 85](#)
2. [Configuring the Egress PE Switch on page 87](#)

### 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-0/0/6.0 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_javae from 208.208.208.208
user@switch# set label-switched-path ip_lsp29_javae 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_javae 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 EX Series Switches on page 29](#)
- [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)
- [Configuring an OSPF Network \(J-Web Procedure\)](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

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

Junos OS MPLS for EX Series switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). 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 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 EX Series Switches on page 29](#)
- [Example: Configuring MPLS-Based Layer 2 VPNs on page 65](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

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 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 89](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)” on page 80](#).

This topic includes:

1. [Configuring CoS on page 92](#)
2. [Configuring an LSP Policer on page 93](#)

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 121.121.121.1/16 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 Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
- [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 on page 94](#)
2. [Configuring an LSP Policer on page 95](#)

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)*.

Related Documentation

- [Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
- [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*

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 \(CLI Procedure\) on page 93](#)
- [Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS \(CLI Procedure\) on page 92](#)
- [Configuring Rewrite Rules for EXP Classifiers on MPLS Networks \(CLI Procedure\)](#)
- [Configuring CoS on Provider Switches of an MPLS Network \(CLI Procedure\) on page 96](#)
- [Defining CoS Rewrite Rules \(CLI Procedure\)](#)

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 103.

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 Switches Using Circuit Cross-Connect \(CLI Procedure\)”](#) on page 89.
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)”](#) on page 80.



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 on page 65](#)
- [Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) on page 101](#)
- [Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches on page 20](#)

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 85.
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)”](#) on page 80.



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 on page 65](#)
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) on page 98](#)

- [Understanding Using MPLS-Based Layer 2 and Layer 3 VPNs on EX Series Switches on page 20](#)

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

You can use 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 Provider Switches \(CLI Procedure\)” on page 80](#).



NOTE: You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

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-name 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 EX Series Switches on page 29](#)
- [Example: Configuring MPLS-Based Layer 2 VPNs on page 65](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

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

You can use configure an 802.1Q VLAN as an MPLS-based Layer 2 circuit 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 Provider Switches \(CLI Procedure\)” on page 80](#).



NOTE: You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

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 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
```

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 EX Series Switches on page 29](#)
- [Example: Configuring MPLS-Based Layer 2 VPNs on page 65](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

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 Provider Switches \(CLI Procedure\)” on page 80](#).



NOTE: You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

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 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 EX Series Switches on page 29](#)
- [Example: Configuring MPLS-Based Layer 2 VPNs on page 65](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

Configuring IPv6 Tunneling for MPLS (CLI Procedure)

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 EX Series Switches on page 29](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

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, liveliness 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 Switches Using Circuit Cross-Connect \(CLI Procedure\)” on page 89](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)” on page 80](#).

This topic describes how to configure an ingress PE switch, one or more provider switches, and an egress PE switch for static LSP:

1. [Configuring the Ingress PE Switch on page 112](#)
2. [Configuring the Provider and the Egress PE Switch on page 113](#)

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 EX Series Switches on page 29](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

Configuring Bidirectional Forwarding Detection for MPLS (CLI Procedure)

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 on page 114](#)
- [Configuring BFD on Provider Edge and Provider Switches for an RSVP-Based LSP on page 116](#)

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 85](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)” on page 80](#).

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 85](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches \(CLI Procedure\)” on page 80](#).

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 EX Series Switches on page 29](#)
- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Understanding Junos OS MPLS Components for EX Series Switches on page 5](#)

CHAPTER 4

Configuration Statements

- [\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches on page 119](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 120](#)
- [\[edit protocols mpls\] Configuration Statement Hierarchy on EX Series Switches on page 121](#)

[\[edit interfaces\] Configuration Statement Hierarchy on EX Series Switches](#)

Each of the following topics lists the statements at a subhierarchy of the **[edit interfaces]** hierarchy:

- [\[edit interfaces ae\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces ge\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces interface-range\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces lo\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces me\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces vlan\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces vme\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit interfaces xe\] Configuration Statement Hierarchy on EX Series Switches](#)

Related Documentation

- [EX Series Switches Interfaces Overview](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring a Layer 3 Subinterface \(CLI Procedure\)](#)
- [Configuring Routed VLAN Interfaces \(CLI Procedure\)](#)
- [Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis \(CLI Procedure\)](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

[\[edit protocols\]](#) Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit protocols]** hierarchy:

- [\[edit protocols bfd\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols bgp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols connections\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols dcbx\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols dot1x\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols igmp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols igmp-snooping\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols isis\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols lacp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols link-management\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols lldp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols lldp-med\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mld\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mld-snooping\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mpls\]](#) Configuration Statement Hierarchy on EX Series Switches on [page 121](#)
- [\[edit protocols msdp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mstp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols mvrp\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols neighbor-discovery\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols oam\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols ospf\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols ospf3\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols pim\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols rip\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols ripng\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols router-advertisement\]](#) Configuration Statement Hierarchy on EX Series Switches
- [\[edit protocols router-discovery\]](#) Configuration Statement Hierarchy on EX Series Switches

- [\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols rsvp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols stp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches](#)

**Related
Documentation**

- [EX Series Switch Software Features Overview](#)
- [Junos® OS for EX Series Switches, Release 12.2](#)

[\[edit protocols mpls\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit protocols mpls]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- [Supported Statements in the \[edit protocols mpls\] Hierarchy Level on page 121](#)
- [Unsupported Statements in the \[edit protocols mpls\] Hierarchy Level on page 122](#)

Supported Statements in the [edit protocols mpls] Hierarchy Level

The following hierarchy shows the **[edit protocols mpls]** configuration statements supported on EX Series switches:

```
protocols {
  mpls {
    class-of-service cos-value;
    disable;
    explicit-null;
    interface (interface-name | all) {
      disable;
    }
    ipv6-tunneling ;
    label-switched-path lsp-name {
      description text-string;
    }
  }
}
```

```

disable;
exclude-slrg;
from address;
ldp-tunneling;
no-cspf;
no-decrement-ttl;
oam {
    bfd-liveness-detection {
        detection-time {
            threshold milliseconds;
        }
        failure-action (make-before-break <teardown-timeout seconds> | teardown);
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        multiplier number;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (1 | automatic);
    }
}
to address;
}
no-cspf;
no-decrement-ttl;
no-propagate-ttl;
static-label-switched-path lsp-name {
    ingress {
        install {
            destination-prefix <active>;
        }
        next-hop (address | interface-name | address/interface-name);
        push out-label;
        to address;
    }
    transit incoming-label {
        description text-string;
        next-hop (address | interface-name | address/interface-name);
        pop;
        swap out-label;
    }
    traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);
}
}

```

Unsupported Statements in the [edit protocols mpls] Hierarchy Level

All statements in the [edit protocols mpls] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as <i>interface-name</i> , are not shown in the statements or hierarchies.	
active	[edit protocols mpls static-label-switched-path ingress install] [edit protocols mpls label-switched-path install]
adaptive	[edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
adjust-interval	[edit protocols mpls label-switched-path auto-bandwidth]
adjust-threshold	[edit protocols mpls label-switched-path auto-bandwidth]
adjust-threshold-overflow-limit	[edit protocols mpls label-switched-path auto-bandwidth]
adjust-threshold-underflow-limit	[edit protocols mpls label-switched-path auto-bandwidth]
admin-down	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
admin-group	[edit protocols mpls] [edit protocols mpls interface] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
admin-group-extended	[edit protocols mpls] [edit protocols mpls interface] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
admin-groups	[edit protocols mpls]
advertisement-hold-time	[edit protocols mpls]
allow-fragmentation	[edit protocols mpls path-mtu]
always-mark-connection-protection-tlv	[edit protocols mpls interface]
associate-backup-pe-groups	[edit protocols mpls label-switched-path]
auto-bandwidth	[edit protocols mpls label-switched-path] [edit protocols mpls statistics]
auto-policing	[edit protocols mpls]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
backup	[edit protocols mpls label-switched-path]
bandwidth	[edit protocols mpls] [edit protocols mpls bandwidth] [edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path bypass] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls static-label-switched-path ingress] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary] [edit protocols mpls static-label-switched-path transit]
bandwidth-model	[edit protocols mpls diffserv-te]
bandwidth-percent	[edit protocols mpls label-switched-path fast-reroute]
bfd-liveness-detection	[edit protocols mpls label-switched-path primary oam] [edit protocols mpls label-switched-path secondary oam] [edit protocols mpls oam]
bypass	[edit protocols mpls static-label-switched-path]
bypass-name	[edit protocols mpls static-label-switched-path ingress link-protection] [edit protocols mpls static-label-switched-path ingress node-protection] [edit protocols mpls static-label-switched-path transit link-protection] [edit protocols mpls static-label-switched-path transit node-protection]
class	[edit protocols mpls auto-policing]
class-of-service	[edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary] [edit protocols mpls static-label-switched-path ingress]
context-identifier	[edit protocols mpls egress-protection]
ct0	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]
ct1	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
ct2	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]
ct3	[edit protocols mpls bandwidth] [edit protocols mpls label-switched-path bandwidth] [edit protocols mpls label-switched-path primary bandwidth] [edit protocols mpls label-switched-path secondary bandwidth]
description	[edit protocols mpls static-label-switched-path bypass] [edit protocols mpls static-label-switched-path ingress]
detection-time	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
diffserv-te	[edit protocols mpls]
drop	[edit protocols mpls auto-policing class]
egress-protection	[edit protocols mpls] [edit protocols mpls label-switched-path]
encoding-type	[edit protocols mpls label-switched-path lsp-attributes]
exclude	[edit protocols mpls admin-group] [edit protocols mpls label-switched-path admin-group] [edit protocols mpls label-switched-path admin-group-extended] [edit protocols mpls label-switched-path primary admin-group] [edit protocols mpls label-switched-path primary admin-group-extended] [edit protocols mpls label-switched-path secondary admin-group] [edit protocols mpls label-switched-path secondary admin-group-extended] [edit protocols mpls label-switched-path fast-reroute]
exclude-slrg	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
expand-loose-hop	[edit protocols mpls]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
failure-action	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
fast-reroute	[edit protocols mpls label-switched-path]
file	[edit protocols mpls label-switched-path primary oam traceoptions] [edit protocols mpls label-switched-path secondary oam traceoptions] [edit protocols mpls label-switched-path traceoptions] [edit protocols mpls statistics] [edit protocols mpls traceoptions]
files	[edit protocols mpls statistics file]
filter	[edit protocols mpls static-label-switched-path ingress policing]
flag	[edit protocols mpls label-switched-path primary oam traceoptions] [edit protocols mpls label-switched-path secondary oam traceoptions] [edit protocols mpls label-switched-path traceoptions] [edit protocols mpls traceoptions]
gpip	[edit protocols mpls label-switched-path lsp-attributes]
hop-limit	[edit protocols mpls] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls label-switched-path]
icmp-tunneling	[edit protocols mpls]
include-all	[edit protocols mpls admin-group] [edit protocols mpls admin-group-extended] [edit protocols mpls label-switched-path admin-group] [edit protocols mpls label-switched-path admin-group-extended] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls label-switched-path primary admin-group] [edit protocols mpls label-switched-path primary admin-group-extended] [edit protocols mpls label-switched-path secondary admin-group] [edit protocols mpls label-switched-path secondary admin-group-extended]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
include-any	[edit protocols mpls admin-group] [edit protocols mpls admin-group-extended] [edit protocols mpls label-switched-path admin-group] [edit protocols mpls label-switched-path admin-group-extended] [edit protocols mpls label-switched-path fast-reroute] [edit protocols mpls label-switched-path primary admin-group] [edit protocols mpls label-switched-path primary admin-group-extended] [edit protocols mpls label-switched-path secondary admin-group] [edit protocols mpls label-switched-path secondary admin-group-extended]
install	[edit protocols mpls label-switched-path]
inter-domain	[edit protocols mpls label-switched-path]
interval	[edit protocols mpls statistics]
least-fill	[edit protocols mpls label-switched-path]
link-protection	[edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress] [edit protocols mpls static-label-switched-path transit]
log-updown	[edit protocols mpls]
loss-priority-high	[edit protocols mpls auto-policing class]
loss-priority-low	[edit protocols mpls auto-policing class]
lsp-attributes	[edit protocols mpls label-switched-path]
make-before-break	[edit protocols mpls label-switched-path secondary oam bfd-liveness-detection failure-action]
maximum-bandwidth	[edit protocols mpls label-switched-path auto-bandwidth]
metric	[edit protocols mpls egress-protection context-identifier] [edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress]
mib-mpls-show-p2mp	[edit protocols mpls]
mimum-bandwidth	[edit protocols mpls label-switched-path auto-bandwidth]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
minimum-interval	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path primary oam bfd-liveness-detection transit-interval] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection transit-interval]
minimum-receive-interval	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
monitor-bandwidth	[edit protocols mpls label-switched-path auto-bandwidth]
most-fill	[edit protocols mpls label-switched-path]
mpls-lsp-traps	[edit protocols mpls log-updown no-trap]
mpls-tp-mode	[edit protocols mpls oam]
mtu-signaling	[edit protocols mpls path-mtu rsvp]
multiplier	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
next-hop	[edit protocols mpls path] [edit protocols mpls static-label-switched-path bypass]
next-next-label	[edit protocols mpls static-label-switched-path ingress node-protection] [edit protocols mpls static-label-switched-path transit node-protection]
no-adaptation	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
no-auto-policing	[edit protocols mpls static-label-switched-path ingress policing]
no-cspf	[edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
no-decrement-ttl	[edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
node-link-protection	[edit protocols mpls label-switched-path]
node-protection	[edit protocols mpls static-label-switched-path ingress] [edit protocols mpls static-label-switched-path transit]
no-exclude	[edit protocols mpls label-switched-path fast-reroute]
no-include-all	[edit protocols mpls label-switched-path fast-reroute]
no-include-any	[edit protocols mpls label-switched-path fast-reroute]
no-install-to-address	[edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress]
no-record	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
no-remote-trace	[edit protocols mpls label-switched-path oam traceoptions] [edit protocols mpls label-switched-path secondary oam traceoptions]
no-syslog	[edit protocols mpls log-updown]
no-trap	[edit protocols mpls log-updown]
no-world-readable	[edit protocols mpls statistics file]
number	[edit protocols mpls auto-policing class]
oam	[edit protocols mpls] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
optimize-aggressive	[edit protocols mpls]
optimize-hold-dead-delay	[edit protocols mpls]
optimize-switchover-delay	[edit protocols mpls]
optimize-timer	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
path	[edit protocols mpls]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
path-mtu	[edit protocols mpls]
policing	[edit protocols mpls static-label-switched-path ingress] [edit protocols mpls label-switched-path]
preference	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls static-label-switched-path ingress] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
primary	[edit protocols mpls egress-protection context-identifier]
priority	[edit protocols mpls] [edit protocols mpls diffserv-te te-class-matrix tex] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
protection-revert-time	[edit protocols mpls interface static]
protector	[edit protocols mpls egress-protection context-identifier]
push	[edit protocols mpls static-label-switched-path bypass]
random	[edit protocols mpls label-switched-path]
record	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
retry-limit	[edit protocols mpls label-switched-path]
retry-timer	[edit protocols mpls label-switched-path]
revert-timer	[edit protocols mpls] [edit protocols mpls label-switched-path]
rfc3812-traps	[edit protocols mpls log-updown no-trap]
rsvp	[edit protocols mpls path-mtu]
rsvp-error-hold-time	[edit protocols mpls]
select	[edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
signal-bandwidth	[edit protocols mpls label-switched-path lsp-attributes]
size	[edit protocols mpls statistics file]
smart-optimize-timer	[edit protocols mpls]
soft-preemption	[edit protocols mpls label-switched-path]
srlg	[edit protocols mpls interface]
standby	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path primary] [edit protocols mpls label-switched-path secondary]
static	[edit protocols mpls interface]
statistics	[edit protocols mpls]
switch-away-lsps	[edit protocols mpls interface]
switching-type	[edit protocols mpls label-switched-path lsp-attributes]
syslog	[edit protocols mpls log-updown]
tex	[edit protocols mpls diffserv-te te-class-matrix]
teardown	[edit protocols mpls label-switched-path secondary oam bfd-liveness-detection failure-action]
te-class-matrix	[edit protocols mpls diffserv-te]
template	[edit protocols mpls label-switched-path]
threshold	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection transmit-interval] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection detection-time] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection transmit-interval]
to	[edit protocols mpls static-label-switched-path bypass]
traceoptions	[edit protocols mpls] [edit protocols mpls label-switched-path] [edit protocols mpls label-switched-path oam] [edit protocols mpls label-switched-path primary oam] [edit protocols mpls label-switched-path secondary oam]

Table 18: Unsupported [edit protocols mpls] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
traffic-class	[edit protocols mpls diffserv-te te-class-matrix tex]
transmit-interval	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
trap	[edit protocols mpls log-updown]
trap-path-down	[edit protocols mpls log-updown]
trap-path-up	[edit protocols mpls log-updown]
version	[edit protocols mpls label-switched-path primary oam bfd-liveness-detection] [edit protocols mpls label-switched-path secondary oam bfd-liveness-detection]
world-readable	[edit protocols mpls statistics file]

**Related
Documentation**

- [Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
- [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
- [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)
- [Junos OS MPLS for EX Series Switches Overview on page 3](#)
- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches](#)

bfd-liveness-detection (Protocols MPLS)

Syntax	<pre> bfd-liveness-detection { failure-action { make-before-break teardown-timeout <i>seconds</i>; teardown; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; minimum-transmit-interval <i>milliseconds</i>; multiplier <i>detection-time-multiplier</i>; } </pre>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i> oam], [edit protocols mpls oam]
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>failure-action option added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 12.2 for EX Series switches.</p>
Description	Enable Bidirectional Forwarding Detection (BFD) for all of the MPLS LSPs or for just a specific LSP.
Options	<p>minimum-interval—Minimum transmit and receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-receive-interval—Minimum receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-transmit-interval—Minimum transmit interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>multiplier—Detection time multiplier. Range: 1 through 255 Default: 3</p> <p>The failure-action statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring BFD for MPLS IPv4 LSPs • Configuring Bidirectional Forwarding Detection for MPLS (CLI Procedure) on page 114

connections (MPLS)

Syntax	<pre>connections { remote-interface-switch <i>connection-name</i> { interface <i>interface-name.unit-number</i>; transmit-lsp <i>label-switched-path</i>; receive-lsp <i>label-switched-path</i>; } }</pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	<p>Define the connection between two circuits in a CCC connection.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring MPLS on EX Series Switches on page 29• <i>Junos OS MPLS Applications Configuration Guide</i>

description (Protocols Layer 2 VPN)

Syntax	<pre>description <i>text</i>;</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i> interface <i>interface-name</i>]
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	<ul style="list-style-type: none">• <i>Configuring the Local Site on PE Routers in Layer 2 VPNs</i>• Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) on page 98

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-bridge ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls ethernet-vpls-fr ether-vpls-over-atm-llc ethernet-vpls-ppp extended-frame-relay-ccc extended-frame-relay-ether-type-tcc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-port-ccc frame-relay-tcc generic-services multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vci-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces rlsq <i>number:number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1 for PTX Series Packet Transport Switches (flexible-ethernet-services , ethernet-ccc , and ethernet-tcc options only).
Description	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
Default	ppp —Use serial PPP encapsulation.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Use ATM PVC encapsulation.</p> <p>cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.</p> <p>cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.</p> <p>cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.</p> <p>ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.</p> <p>ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.</p> <p>ethernet-over-atm—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination</p>

IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

ethernet-vpls-ppp—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

ether-vpls-over-atm-llc—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-bridge—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

generic-services—Use generic services encapsulation for services with a hierarchical scheduler.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



.....
NOTE: Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
.....

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

**Related
Documentation**

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- *Configuring Layer 2 Switching Cross-Connects Using CCC*
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Switches*
- [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) on page 98](#)
- *Configuring MPLS LSP Tunnel Cross-Connects Using CCC*
- *Configuring TCC*
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- *Defining the Encapsulation for Switching Cross-Connects*
- *Understanding Encapsulation on an Interface*

encapsulation-type (Layer 2 VPNs)

Syntax	<code>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);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls neighbor <i>address</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p>
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	<p>atm-aal5—ATM Adaptation Layer (AAL/5)</p> <p>atm-cell—ATM cell relay</p> <p>atm-cell-port-mode—ATM cell relay port promiscuous mode</p> <p>atm-cell-vc-mode—ATM VC cell relay nonpromiscuous mode</p> <p>atm-cell-vp-mode—ATM virtual path (VP) cell relay promiscuous mode</p> <p>cesop—CESOP-based Layer 2 VPN</p> <p>cisco-hdlc—Cisco Systems-compatible HDLC</p> <p>ethernet—Ethernet</p> <p>ethernet-vlan—Ethernet VLAN</p> <p>frame-relay—Frame Relay</p> <p>frame-relay-port-mode—Frame Relay port mode</p> <p>interworking—Layer 2.5 interworking VPN</p> <p>ppp—PPP</p> <p>satop-e1—SATSOP-E1-based Layer 2 VPN</p>

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

Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Local Site on PE Routers in Layer 2 VPNs</i>• <i>Configuring VPLS Routing Instances</i>• <i>Configuring Interfaces for Layer 2 Circuits</i>• Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) on page 98
------------------------------	---

exp

Syntax	<pre>exp classifier-name { import (classifier-name default); forwarding-class class-name { loss-priority level { code-points [aliases] [3-bit-patterns]; } } }</pre>
Hierarchy Level	[edit class-of-service classifiers], [edit class-of-service code-point-aliases], [edit class-of-service interfaces <i>interface-name</i> unit <i>logical-unit-number</i> rewrite-rules], [edit class-of-service rewrite-rules]
Release Information	Statement introduced in Junos OS Release 10.1 for EX Series switches.
Description	<p>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.</p> <p>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.</p>
Options	<p>classifier-name—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding Using CoS with MPLS Networks on EX Series Switches on page 12• Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 89• Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 85• Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) on page 96

fec

Syntax	<pre> fec <i>fec-address</i> { bfd-liveness-detection { detection-time threshold <i>milliseconds</i>; ecmp; failure-action { remove-nexthop; remove-route; } holddown-interval <i>milliseconds</i>; ingress-policy <i>ingress-policy-name</i>; minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; minimum-transmit-interval <i>milliseconds</i>; multiplier <i>detection-time-multiplier</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (0 1 automatic); } no-bfd-liveness-detection; periodic-traceroute { disable; exp <i>exp-value</i>; fanout <i>fanout-value</i>; frequency <i>minutes</i>; paths <i>number-of-paths</i>; retries <i>retry-attempts</i>; source <i>address</i>; ttl <i>ttl-value</i>; wait <i>seconds</i>; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-systems-name</i> 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.
Description	Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).
Options	<p><i>fec-address</i>—Specify the FEC address.</p> <p>The other statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Configuring BFD for LDP LSPs*

instance-type

Syntax	<code>instance-type type;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
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.
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:

- **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. For more information about configuring a **layer2-control** instance type, see the *Junos OS Layer 2 Configuration Guide*.
- **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 only) Provide support for Layer 2 bridging. Use this routing instances type to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and separates its VLAN identifier space. For more information about configuring a virtual switch instance type, see the *Junos OS Layer 2 Configuration Guide*, and the *JUNOS® MX Series 3D Universal Edge Routers Solutions, Release 12.3*.
- **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

- *Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches*
- *Configuring Routing Instances on PE Routers in VPNs*
- *Configuring Virtual Routing Instances (CLI Procedure)*
- *Configuring Virtual Router Routing Instances*
- *Example: Configuring Filter-Based Forwarding on the Source Address*
- *Example: Configuring Filter-Based Forwarding on Logical Systems*
- *Junos OS Layer 2 Configuration Guide*
- *JUNOS® MX Series 3D Universal Edge Routers Solutions, Release 12.3*

interface (MPLS)

Syntax	interface (all <i>interface-name</i>);
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	<p>all—All interfaces on the switch.</p> <p><i>interface-name</i>—Name of an interface:</p> <ul style="list-style-type: none"> • Aggregated Ethernet—aex • Gigabit Ethernet—ge-fpc/pic/port
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring MPLS on EX Series Switches on page 29 • Configuring MPLS on Provider Edge Switches (CLI Procedure) • Configuring MPLS on Provider Switches (CLI Procedure) on page 80

l2circuit

Syntax	<pre> l2circuit { local-switching { interface <i>interface-name</i> { description <i>text</i>; } end-interface { interface <i>interface-name</i>; protect-interface <i>interface-name</i>; } ignore-mtu-mismatch; protect-interface <i>interface-name</i>; } } neighbor <i>address</i> { interface <i>interface-name</i> { bandwidth (<i>bandwidth</i> <i>ctnumber bandwidth</i>); community <i>community-name</i>; (control-word no-control-word); description <i>text</i>; encapsulation-type <i>type</i>; ignore-encapsulation-mismatch; ignore-mtu-mismatch; mtu <i>mtu-number</i>; protect-interface <i>interface-name</i>; pseudowire-status-tlv; psn-tunnel-endpoint <i>address</i>; virtual-circuit-id <i>identifier</i>; } } traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	<p>Enables a Layer 2 circuit.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring 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*

l2vpn

Syntax	<pre> l2vpn { (control-word no-control-word); encapsulation-type type; traceoptions { file filename <files number> <size size> <world-readable no-world-readable>; flag flag <flag-modifier> <disable>; } site site-name { site-identifier identifier; site-preference preference-value { backup; primary; } interface interface-name { description text; remote-site-id remote-site-id; } } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	<p>Enable a Layer 2 VPN routing instance on a PE router or switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Local Site on PE Routers in Layer 2 VPNs</i> • Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) on page 98

label-switched-path

Syntax	label-switched-path <i>lsp-name</i> to <i>remote-provider-edge-switch</i> ;
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	<p><i>lsp-name</i> —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.</p> <p><i>remote-provider-edge-switch</i> —Either the loopback address or the switch address.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring MPLS on EX Series Switches on page 29• Configuring MPLS on Provider Edge Switches (CLI Procedure)• Junos OS MPLS Applications Configuration Guide

ldp

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

```

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

```

```

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

```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches.
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.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Minimum LDP Configuration</i> • <i>Enabling and Disabling LDP</i>

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.
Description	<p>Enable MPLS on the switch.</p> <p>The remaining statements are explained separately.</p>
Default	MPLS is disabled.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring MPLS on EX Series Switches on page 29 • Configuring MPLS on Provider Edge Switches (CLI Procedure) • Configuring MPLS on Provider Switches (CLI Procedure) on page 80 • Junos OS MPLS Applications Configuration Guide

neighbor (Protocols Layer 2 Circuit)

Syntax	<pre> neighbor address { interface interface-name { backup-neighbor address { community name; psn-tunnel-endpoint address; standby; static; virtual-circuit-id number; } bandwidth (bandwidth ctnumber bandwidth); community community-name; (control-word no-control-word); description text; ignore-encapsulation-mismatch; ignore-mtu-mismatch; mtu mtu-number; no-revert; protect-interface interface-name; pseudowire-status-tlv; psn-tunnel-endpoint address; revert-time seconds; static { incoming-label label; outgoing-label label; } switchover-delay milliseconds; virtual-circuit-id identifier; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> 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	<p>address—IP address of a neighboring router or switch.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

Related Documentation • [Configuring Interfaces for Layer 2 Circuits](#)

path

Syntax	<pre>path destination { <address hostname> <strict loose> }</pre>
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	<p>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.</p> <p>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.</p> <p>Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.</p> <p>hostname —(Optional) See address.</p> <p>Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.</p> <p>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.</p> <p>Default: strict</p> <p>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.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	• Configuring Path Protection in an MPLS Network (CLI Procedure) on page 81

periodic-traceroute

Syntax	<pre>periodic-traceroute { disable; exp <i>exp-value</i>; fanout <i>fanout-value</i>; frequency <i>minutes</i>; paths <i>number-of-paths</i>; retries <i>retry-attempts</i>; source <i>address</i>; ttl <i>ttl-value</i>; wait <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i>], [edit protocols ldp oam], [edit protocols ldp oam fec <i>fec-address</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Support added at the [edit protocols ldp oam] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam] hierarchy levels in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.2 for EX Series switches.
Description	Enable tracing of forwarding equivalence classes (FECs) for LDP LSPs.
Options	<p>disable—(Optional) Disable tracing for a specific FEC. This option is available at the [edit protocols ldp oam fec <i>fec-address</i> periodic-traceroute] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i> periodic-traceroute] hierarchy levels only.</p> <p>exp <i>exp-value</i>—(Optional) Specify the class of service to use when sending probes. Default: 7 Range: 0 through 7</p> <p>fanout <i>fanout-value</i>—(Optional) Specify the maximum number of next hops to search per node. Default: 16 Range: 1 through 16</p> <p>frequency <i>minutes</i>—(Optional) Specify the interval between traceroute attempts. Default: 60 minutes Range: 15 through 120 minutes</p> <p>paths <i>number-of-paths</i>—(Optional) Specify the maximum number of paths to search. Default: 3 Range: 1 through 255</p>

retries *retry-attempts*—(Optional) Specify the number of attempts to send a probe to a specific node before giving up.

Default: 3

Range: 1 through 9

source address—(Optional) Specify the IPv4 source address to use when sending probes.

ttl value—(Optional) Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.

Default: 64

Range: 1 through 255

wait seconds—(Optional) Specify the wait interval before resending a probe packet.

Default: 10 seconds

Range: 5 though 15 seconds

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

Related Documentation	<ul style="list-style-type: none">• <i>Configuring LDP LSP Traceroute</i>
------------------------------	---

policing

Syntax	<code>policing (filter <i>filter-name</i> no-automatic-policing);</code>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i>] [edit interfaces <i>interface-id</i> unit <i>number-of-logical-unit</i> family inet address <i>ip-address</i>]
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: <ul style="list-style-type: none">• To the LSP for MPLS over CCC.• To the customer-edge interface for IP over MPLS.
Options	filter <i>filter-name</i> —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	<ul style="list-style-type: none">• <i>policer</i>• <i>Configuring Policers to Control Traffic Rates (CLI Procedure)</i>• Configuring CoS on an MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) on page 93• Configuring CoS on an MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure) on page 92

primary

Syntax	<code>primary <i>path-name</i>;</code>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i>]
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	<i>path-name</i> —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	<ul style="list-style-type: none">• Configuring Path Protection in an MPLS Network (CLI Procedure) on page 81

remote-interface-switch

Syntax	<pre>remote-interface-switch <i>connection-name</i> { interface <i>interface-name.unit-number</i>; receive-lsp <i>label-switched-path</i>; transmit-lsp <i>label-switched-path</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> 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	<p><i>connection-name</i>—Connection name.</p> <p><i>interface interface-name.unit-number</i>—Interface name. Include the logical portion of the name, which corresponds to the logical unit number of the CCC interface.</p> <p><i>receive-lsp label-switched-path</i>—Name of the LSP from the connection's source. This LSP name was specified by the <i>label-switched-path</i> statement on the remote provider edge switch in the protocols mpls stanza.</p> <p><i>transmit-lsp label-switched-path</i>—Name of the LSP to the connection's destination. This LSP name was specified by the <i>label-switched-path</i> statement on the local provider edge switch in the protocols mpls stanza.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring MPLS LSP Tunnel Cross-Connects Using CCC • Example: Configuring MPLS on EX Series Switches on page 29 • Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 89 • Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 85 • Junos OS MPLS Applications Configuration Guide


remote-site-id

Syntax	<code>remote-site-id remote-site-ID;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i> interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i> interface <i>interface-name</i>]
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	<i>remote-site-ID</i> —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	<ul style="list-style-type: none">• Configuring the Local Site on PE Routers in Layer 2 VPNs• Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) on page 98

revert-timer

Syntax	<code>revert-timer <i>seconds</i>;</code>
Hierarchy Level	[edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	<p>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.</p> <p>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.</p>
Default	60 seconds
Options	<p><i>seconds</i> —Value in seconds.</p> <p>Range: 0 through 65,535 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Path Protection in an MPLS Network (CLI Procedure) on page 81

route-distinguisher

Syntax	<code>route-distinguisher (as-number:id ip-address:id);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>], [edit routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Support at [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>] hierarchy level introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p>
Description	<p>Specify an identifier attached to a route, enabling you to distinguish to which VPN or VPLS the route belongs. Each routing instance must have a unique route distinguisher associated with it. The route distinguisher 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.</p> <p>For Layer 2 VPNs and VPLS, if you configure the l2vpn-use-bgp-rules statement, you must configure a unique route distinguisher for each PE router participating in the routing instance.</p> <p>For other types of VPNs, we recommend that you use a unique route distinguisher for each PE router participating in specific routing instance. Although you can use the same route distinguisher on all PE routers for the same VPN routing instance, if you use a unique route distinguisher, you can determine the CE router from which a route originated within the VPN.</p>
	<div>  <p>CAUTION: We strongly recommend that if you change a route distinguisher that has already been configured, make the change during a maintenance window, as follows:</p> <ol style="list-style-type: none"> 1. Deactivate the routing instance. 2. Change the route distinguisher. 3. Activate the routing instance. <p>This is not required if you are configuring the route distinguisher for the first time.</p> </div>
Options	<p>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</p>

4-byte value, the administrative number is a 2-byte value. A route distinguisher consisting of a 4-byte AS number and a 2-byte administrative number is defined as a type 2 route distinguisher in RFC 4364 *BGP/MPLS IP Virtual Private Networks (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 a route distinguisher that includes a 4-byte AS number, append the letter “L” to the end of the AS number. For example, a route distinguisher with the 4-byte AS number 7,765,000 and an administrative number of 1,000 is represented as 7765000L: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 *Establishing a Peer Relationship Between a 4-Byte Capable Router and a 2-Byte Capable Router Using a 4-Byte AS Number* in the *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*.

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 Route Target Filtering for VPNs*
 - *Configuring Routing Instances on PE Routers in VPNs*
 - [Configuring an MPLS-Based Layer 2 VPN \(CLI Procedure\) on page 98](#)
 - [Configuring an MPLS-Based Layer 3 VPN \(CLI Procedure\) on page 101](#)
 - *Understanding 4-Byte AS Numbers and Route Distinguishers* in the [Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview](#)
 - *l2vpn-use-bgp-rules*

rsvp

Syntax	rsvp;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	<p>Enable Resource Reservation Protocol (RSVP) signaling.</p> <p>You must include the rsvp statement in the configuration to enable RSVP on the router.</p> <p>The primary purpose of RSVP in Junos OS for EX Series switches is to support dynamic signaling within label switched paths (LSPs).</p>
Default	RSVP is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Minimum RSVP Configuration</i>• Example: Configuring MPLS on EX Series Switches on page 29• Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 89• Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 85• Configuring MPLS on Provider Switches (CLI Procedure) on page 80• <i>Junos OS MPLS Applications Configuration Guide</i>

secondary

Syntax	<code>secondary <i>path-name</i> { standby; }</code>
Hierarchy Level	[edit protocols <code>mpls label-switched-path</code> <i>lsp-name</i>]
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	<i>path-name</i> —Name of a secondary path that you created with the path statement. The remaining statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Path Protection in an MPLS Network (CLI Procedure) on page 81

signaling

Syntax	signaling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet-mdt], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet-mvpn], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet-mdt], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet-mvpn], [edit routing-instances <i>routing-instance-name</i> protocols bgp family inet-mdt], [edit routing-instances <i>routing-instance-name</i> protocols bgp family inet-mvpn], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet-mdt], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> 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	<ul style="list-style-type: none">• <i>Example: Configuring Source-Specific Multicast for Draft-Rosen Multicast VPNs</i>• Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) on page 98

site (Layer 2 Circuits)

Syntax	<pre> site <i>site-name</i> { site-identifier <i>identifier</i>; site-preference <i>preference-value</i> { backup; primary; } interface <i>interface-name</i> { description <i>text</i>; remote-site-id <i>remote-site-ID</i>; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn]
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 site name, site identifier, and interfaces connecting to the site. Allows you to configure a remote site ID for remote sites.
Options	<p>site-identifier <i>identifier</i>—Numerical identifier for the site used as a default reference for the remote site ID.</p> <p><i>site-name</i>—Name of the site.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Local Site on PE Routers in Layer 2 VPNs • Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) on page 98

site-identifier (Layer 2 Circuits)

Syntax	site-identifier <i>identifier</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols l2vpn site <i>site-name</i>]
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	<i>identifier</i> —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	<ul style="list-style-type: none">• Configuring the Local Site on PE Routers in Layer 2 VPNs• Configuring an MPLS-Based Layer 2 VPN (CLI Procedure) on page 98

standby

Syntax	standby;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	Enable the path to remain up at all times to provide instant switchover if connectivity problems occur.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Hot Standby of Secondary Paths• Configuring Path Protection in an MPLS Network (CLI Procedure) on page 81

traffic-engineering

Syntax	traffic-engineering;
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.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring MPLS on EX Series Switches on page 29 • Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 89 • Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 85 • Configuring MPLS on Provider Switches (CLI Procedure) on page 80 • Configuring an OSPF Network (J-Web Procedure) • Junos OS MPLS Applications Configuration Guide

vrf-table-label

Syntax	vrf-table-label;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Map the inner label of a packet to a specific VPN routing and forwarding (VRF) table. This allows the examination of the encapsulated IP header. All routes in the VRF configured with this option are advertised with the label allocated per VRF.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Filtering 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

vrf-target

Syntax	<pre>vrf-target { community; import community-name; export community-name; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>] [edit routing-instances <i>routing-instance-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols vpls mesh-group <i>mesh-group-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Specify a 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.</p> <p>You can still create more complex policies by explicitly configuring VRF import and export policies using the import and export options.</p>
Options	<p>community—Community name.</p> <p>import community-name—Allowed communities accepted from neighbors.</p> <p>export community-name—Allowed communities sent to neighbors.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">Configuring Policies for the VRF Table on PE Routers in VPNs

PART 3

Administration

- [Routine Monitoring on page 175](#)
- [Operational Commands on page 183](#)

CHAPTER 5

Routine Monitoring

- [Verifying That MPLS Is Working Correctly on page 175](#)
- [Verifying Path Protection in an MPLS Network on page 178](#)

Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly on EX Series switches, perform the following tasks:

1. [Verifying the Physical Layer on the Switches on page 175](#)
2. [Verifying the Routing Protocol on page 176](#)
3. [Verifying the Core Interfaces Being Used for the MPLS Traffic on page 176](#)
4. [Verifying RSVP on page 176](#)
5. [Verifying the Assignment of Interfaces for MPLS Label Operations on page 177](#)
6. [Verifying the Status of the CCC on page 177](#)
7. [Verifying traceroute for Layer 3 VPN on page 178](#)

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 ge- terse**

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up			
ge-0/0/1.0	up	up	ccc		
ge-0/0/2.0	up	up	ccc		
ge-0/0/3.0	up	up	eth-switch		
ge-0/0/4.0	up	up	eth-switch		
ge-0/0/5.0	up	up	inet	10.1.5.1/24	
mpls					
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 interfaces **ge-0/0/1.0** and **ge-0/0/2.0** are configured as circuit cross-connect. The Local and Remote columns do not display IP addresses, because the **inet family** is not configured for CCC interfaces. The output for the protocol family of the core interfaces (**ge-0/0/0.5** and **ge-0/0/0.6**), 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. If you have configured IS-IS as the routing protocol, use the **show isis adjacency** 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
127.1.1.2	ge-0/0/5	Full	10.10.10.10	128	39

Meaning The **show ospf neighbor** command displays the status of the routing protocol 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-0/05	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**.

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. You should perform this task only on the provider edge switches.

Action user@switch> `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  0 127.1.2.1         Push 299792 540 2 ge-0/0/5.0
```

Meaning This output shows that CCC has been set up on interface **ge-0/0/1.0**. The switch receives ingress traffic on **ge-0/0/1.0** with label **299776**. It pops that label and swaps it to label **299792**, which it pushes out on interface **ge-0/0/5.0**.

Verifying the Status of the CCC

Purpose Verify the status of the CCC. You should perform this task only on the provider edge switches.

```

Action      user@switch> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)
UN -- uninitialized
NP -- not present
WE -- wrong encapsulation
DS -- disabled
Dn -- down
-> -- only outbound conn is up
<- -- only inbound conn is up
Up -- operational
RmtDn -- remote CCC down
Restart -- restarting

Legend for connection types
if-sw: interface switching
rmt-if: remote interface switching
lsp-sw: LSP switching
tx-p2mp-sw: transmit P2MP switching
rx-p2mp-sw: receive P2MP switching

Legend for circuit types
intf -- interface
tlsp -- transmit LSP
rlsp -- receive LSP

Connection/Circuit      Type      St      Time last up      # Up trans
ge1-to-pe2              rmt-if    Up      Feb 17 05:00:09    1
  ge-0/0/1.0             intf      Up
    lsp_to_pe1_ge1       tlsp      Up
    lsp_to_pe2_ge1       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**.

Verifying traceroute for Layer 3 VPN

Purpose Verify the route that packets take to a specified network host..

```

Action    user@switch> run traceroute routing-instance vpn1 20.0.1.1
           traceroute to 20.0.1.1 (20.0.1.1), 30 hops max, 40 byte packets
           1  * * *
           2  * * *
           3  * 20.0.1.1 (20.0.1.1) 6.077 ms 5.513 ms

```

Meaning The **run traceroute** command displays route that packets take to a specified network host. You can also use traceroute as a debugging tool to locate points of failure in a network. This output verifies that the Layer 3 VPN named *vpn1* is correctly configured and that the packets are able to traverse across *vpn1* to the destination host *20.0.1.1*.

Related Documentation

- [Configuring MPLS on Provider Edge Switches \(CLI Procedure\)](#)
- [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)

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 on page 179](#)
2. [Verifying the RSVP-Enabled Interfaces on page 179](#)
3. [Verifying a Secondary Path on page 180](#)

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
      Active Subscr- Static      Available  Reserved  Highwater
Interface State resv  iption BW      BW      BW      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 \(CLI Procedure\) on page 81](#)
 - [Understanding MPLS and Path Protection on EX Series Switches on page 11](#)

CHAPTER 6

Operational Commands

clear mpls lsp

Syntax	<pre>clear mpls lsp <autobandwidth> <logical-system (all <i>logical-system-name</i>)> <name <i>name</i>> <optimize optimize-aggressive> <path <i>regular-expression</i>> <statistics></pre>
Syntax (EX Series Switches)	<pre>clear mpls lsp <autobandwidth> <name <i>name</i>> <optimize optimize-aggressive> <path <i>regular-expression</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
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	<p>none—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p>autobandwidth—(Optional) Clear LSP autobandwidth counters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>name</i>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p>optimize optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p> <p>path <i>regular-expression</i>—(Optional) Clear the specific LSP path matching the specified regular expression.</p> <p>statistics—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (name and path options) on transit routers.</p>
----------------	--

Required Privilege Level clear

Related Documentation

- [show mpls lsp on page 243](#)
- [show rsvp session on page 272](#)

List of Sample Output [clear mpls lsp on page 185](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

`clear mpls lsp`

```
user@host> clear mpls lsp
```

clear rsvp session

Syntax	<pre>clear rsvp session <connection-destination address> <connection-source address> <gracefully> <logical-system (all logical-system-name)> <lsp-id identifier> <name name> <optimize-fast-reroute> <tunnel-id identifier></pre>
Syntax (EX Series Switches)	<pre>clear rsvp session <connection-destination address> <connection-source address> <gracefully> <lsp-id identifier> <name name> <optimize-fast-reroute> <tunnel-id identifier></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Reset and restart Resource Reservation Protocol (RSVP) sessions.
Options	<p>none—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.</p> <p>connection-source address—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p>connection-destination address—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p>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.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>lsp-id identifier—(Optional) LSP identifier (source port) for the RSVP sender template.</p> <p>name name—(Optional) Reset and restart the specified RSVP session.</p> <p>optimize-fast-reroute—(Optional) Begin fast reroute optimization.</p> <p>tunnel-id identifier—(Optional) Tunnel identifier (destination port) for the RSVP session.</p>

Required Privilege Level clear

Related Documentation

- [clear mpls lsp on page 184](#)
- [show rsvp session on page 272](#)

List of Sample Output [clear rsvp session on page 187](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear rsvp session](#)

```
user@host> clear rsvp session
```

clear rsvp statistics

Syntax	clear rsvp statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear rsvp statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Clear Resource Reservation Protocol (RSVP) packet and error statistics.
Options	none —Clear RSVP packet and error statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show rsvp statistics on page 281
List of Sample Output	clear rsvp statistics on page 188
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
```

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.
 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.

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).

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

Required Privilege Level network

List of Sample Output [ping mpls l2circuit interface on page 190](#)
[ping mpls l2circuit virtual-circuit detail on page 190](#)
[ping mpls l2circuit interface <interface-name> reply-mode on page 191](#)

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*>
<fec129>
<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.
size and **sweep** options introduced in Junos OS Release 9.6.
reply-mode option and its suboptions introduced in Junos OS Release 10.4R1.
fec129 option introduced in Junos OS Release 12.2.

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.

fec129—(Optional) Ping the LSP for an FEC 129 Layer 2 VPN connection.

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

Required Privilege Level network

List of Sample Output [ping mpls l2vpn instance on page 193](#)
[ping mpls l2vpn instance detail on page 194](#)
[ping mpls l2vpn interface <interface-name> reply-mode on page 194](#)

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 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	<pre>ping mpls l3vpn prefix <i>prefix-name</i> <l3vpn-name> <bottom-label-ttl> <count <i>count</i>> <destination <i>address</i>> <detail> <exp <i>forwarding-class</i>> <logical-system (all <i>logical-system-name</i>)> <size <i>bytes</i>> <source <i>source-address</i>> <sweep></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The size and sweep options were introduced in Junos OS Release 9.6.</p>
Description	<p>Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a ping mpls l3vpn command.</p>
Options	<p>bottom-label-ttl—(Optional) Display the time-to-live value for the bottom label in the label stack.</p> <p>count <i>count</i>—(Optional) Number of ping requests to send. If count is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.</p> <p>destination <i>address</i>—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.</p> <p>detail—(Optional) Display detailed information about the echo requests sent and received.</p> <p>exp <i>forwarding-class</i>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p>l3vpn-name—(Optional) Layer 3 VPN name.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p>prefix <i>prefix-name</i>—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.</p> <p>size <i>bytes</i>—(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.</p>

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

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

Additional Information You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response 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.

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 196](#)
[ping mpls l3vpn detail on page 196](#)

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. When an echo reply is received with an error code, the packets are not counted in the received packets count, and are counted separately..

Sample Output

ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!!!
--- 1sping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
```

```
--- lsping statistics ---  
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls ldp

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

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

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

Additional Information 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 *Junos OS MPLS Applications Configuration 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.

Required Privilege Level network

List of Sample Output [ping mpls ldp fec count on page 199](#)
[ping mpls ldp p2mp root-addr lsp-id on page 199](#)

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--- 1sping 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 lsp-end-point

Syntax	<pre>ping mpls lsp-end-point <i>prefix-name</i> <count <i>count</i>> <destination <i>address</i>> <detail> <exp <i>forwarding-class</i>> <instance <i>routing-instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <size <i>bytes</i>> <source <i>source-address</i>> <sweep></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The size and sweep options were introduced in Junos OS Release 9.6.</p> <p>The instance option was introduced in Junos OS Release 10.0.</p>
Description	<p>Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a ping mpls command.</p>
Options	<p>count <i>count</i>—(Optional) Number of ping requests to send. If count is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.</p> <p>destination <i>address</i>—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.</p> <p>detail—(Optional) Display detailed information about the echo requests sent and received.</p> <p>exp <i>forwarding-class</i>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p>instance <i>routing-instance-name</i>—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p>prefix-name—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.</p> <p>size <i>bytes</i>—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is 88 bytes. If the endpoint is RSVP-based, the minimum size of the packet is 100 bytes. The maximum size in either case is 65468 bytes.</p> <p>source <i>source-address</i>—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).</p> <p>sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).</p>

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

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

Required Privilege Level network

List of Sample Output [ping mpls lsp-end-point detail on page 202](#)

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 rsvp

Syntax ping mpls rsvp
 <lsp-name>
 <count count>
 <destination address>
 <detail>
 <dynamic-bypass>
 <egress egress-address>
 <exp forwarding-class>
 <interface interface-name>
 <logical-system (all | logical-system-name)>
 <manual-bypass>
 <multipoint>
 <size bytes>
 <source source-address>
 <standby standby-path-name>
 <sweep>

Release Information Command introduced before Junos OS Release 7.4.
 The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.
 The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.

Description Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options **count count**—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination address—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.



NOTE: When using the **detail** option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.

In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the **detail** option (even when synchronized using NTP).

dynamic-bypass—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

egress *egress-address*—(Optional) Only the specified egress router or switch responds to the ping request.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

interface—(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the **manual-bypass** option.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

lsp-name—Ping an RSVP-signaled LSP using an LSP name.

manual-bypass—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the **interface** option.

multipoint—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

size *bytes*—(Optional) Size of the LSP ping request packet (100 through 65468 bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

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

standby *standby-path-name*—(Optional) Name of the standby path.

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

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

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

Required Privilege Level network

List of Sample Output [ping mpls rsvp \(Echo Reply Received\) on page 205](#)
[ping mpls rsvp \(Echo Reply with Error Code\) on page 205](#)

[ping mpls rsvp detail on page 205](#)

[ping mpls rsvp multipoint egress detail count on page 205](#)

[ping mpls rsvp multipoint detail count on page 205](#)

[ping mpls rsvp destination detail count size on page 206](#)

[ping mpls rsvp destination detail sweep size on page 206](#)

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

Sample Output

ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
loss
```

ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
Local transmit time: 1205310695s 215737us
Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 ms
Local transmit time: 1205310615s 347317us
Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
```

```

Local transmit time: 1205310615s 347262us
Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
Local transmit time: 1205310615s 347167us
Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

```

ping mpls rsvp destination detail count size

```

user@host> ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
    Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
    Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

```

ping mpls rsvp destination detail sweep size

```

user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
    Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
    Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
    Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
    Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
    Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
    Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
    Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
    Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
    Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
    Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms

```



```
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
    Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
    Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
    Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
    Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
    Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
    Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
    Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
    Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

request mpls lsp adjust-autobandwidth

Syntax	<code>request mpls lsp adjust-autobandwidth</code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><name <i>lsp-name</i>></code>
Syntax (EX Series Switches)	<code>request mpls lsp adjust-autobandwidth</code> <code><name <i>lsp-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	<p>Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).</p> <p>Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.</p> <p>During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as <i>make-before-break</i>.</p>
Options	<p>none—Manually trigger a bandwidth allocation adjustment for all active LSP paths.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>lsp-name</i>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.</p>
Additional Information	<p>For this command to work properly, the following conditions must exist:</p> <ul style="list-style-type: none">• Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the request mpls lsp adjust-autobandwidth command.• The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none">• <i>auto-bandwidth</i>• <i>Configuring Automatic Bandwidth Allocation for LSPs</i>
List of Sample Output	request mpls lsp adjust-auto-bandwidth on page 209
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

Syntax	<pre>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 <i>logical-system-name</i>)> <name> <status></pre>
Syntax (EX Series Switches)	<pre>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></pre>
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	<p>none—Display the standard level of output for all configured CCC connections.</p> <p>all—(Optional) Display all connections.</p> <p>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.</p> <p>interface-switch—(Optional) Display interface switch connections only.</p> <p>lsp-switch—(Optional) Display LSP switch connections only.</p> <p>p2mp-receive-switch—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.</p> <p>p2mp-transmit-switch—(Optional) Display local interface to point-to-multipoint LSP switch connections only.</p> <p>remote-interface-switch—(Optional) Display remote interface switch connections only.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both kinds of connections.</p> <p>history—(Optional) Display information about connection history.</p>

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 19 on page 211](#) describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 19: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On Off]	Whether link monitoring is enabled: On or Off .
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.
Legend for connection types	Type of connection: <ul style="list-style-type: none"> if-sw—Layer 2 switching cross-connect. rmt-if—Remote interface switch. While graceful restart is in progress, rmt-if will display a state (St) of Restart. lsp-sw—LSP stitching cross-connect. While graceful restart is in progress, lsp-sw will display a state (St) of Restart.
Legend for circuit types	Type of circuits: <ul style="list-style-type: none"> intf—Interface circuit. tlsp—Transmit LSP circuit. rlsp—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the Up (operational) state.
# 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)
  UN -- uninitialized
  NP -- not present
  WE -- wrong encapsulation
  DS -- disabled
  Dn -- down
  -> -- only outbound conn is up
  <- -- only inbound conn is up
  Up -- operational
  RmtDn -- remote CCC down
  Restart -- restarting

  Legend for connection types
  if-sw: interface switching
  rmt-if: remote interface switching
  lsp-sw: LSP switching

  Legend for circuit types
  intf -- interface
  tlsp -- transmit LSP
  rlsp -- receive LSP

CCC Graceful restart : Restarting

Connection/Circuit      Type   St      Time last up    # Up trans
IFSW-ed                 if-sw  Up       Aug  5 15:39:15      1
  so-1/0/2.0             intf   Up
  t1-0/1/2.0             intf   Up
SW-db                   rmt-if Restart      0
  so-1/0/3.0             intf   Up
  pro4-ca                tlsp   Dn
  pro4-ac                rlsp   NP
```

show link-management

Syntax	show link-management
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management peer on page 217 • show link-management routing on page 219 • show link-management statistics on page 222 • show link-management te-link on page 224
List of Sample Output	show link-management on page 216
Output Fields	Table 20 on page 213 describes the output fields for the show link-management command. Output fields are listed in the approximate order in which they appear.

Table 20: show link-management Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down .
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.
State	State of the control channel: Up or Down .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295.

Table 20: show link-management Output Fields (*continued*)

Field Name	Field Description
Flags	Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.
TE link name	Name of the traffic-engineered link.
State	State of the traffic-engineered link: Up , Down , or Init .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH .
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
Maximum bandwidth	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
Name	Name of the interface.
State	State of the interface: Up or Down .
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: Yes or No .

Table 20: show link-management Output Fields (*continued*)

Field Name	Field Description
LSP-name	LSP name.

Sample Output

show link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
State: Up, Control address: 10.255.245.4
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    24547      24547 Up          1027      1026
TE links:
  pro4-ba

TE link name: pro4-ba, State: Init
Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
PSC-1,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps,
Available bandwidth: 155.52Mbps
  Name          State Local ID Remote ID      Bandwidth Used  LSP-name
  so-1/0/2      Up          21271      0      155.52Mbps    No
```

show link-management peer

Syntax	show link-management peer <name <i>peer-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display Multiprotocol Label Switching (MPLS) peer link information.
Options	none —Display all peer link information. name <i>peer-name</i> —(Optional) Display information for the specified peer only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 213 • show link-management routing on page 219 • show link-management statistics on page 222 • show link-management te-link on page 224
List of Sample Output	show link-management peer on page 218
Output Fields	Table 21 on page 217 describes the output fields for the show link-management peer command. Output fields are listed in the approximate order in which they appear.

Table 21: show link-management peer Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down .
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.

Table 21: show link-management peer Output Fields (*continued*)

Field Name	Field Description
State	State of the control channel: Up or Down .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295 .
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295 .
Flags	Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.

Sample Output

show link-management peer

```

user@host> show link-management peer
Peer name: sonet, System identifier: 41448
State: Up, Control address: 70.70.70.70
Hello interval: 10000, Hello dead interval: 30000
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    3265           0 ConfSnd         1          0 R
TE links:
  to-sonet

```

show link-management routing

Syntax	show link-management routing <peer <name <i>name</i> > te-link <name <i>name</i> >> <resource <name <i>name</i> >>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.
Options	<p>none—Display all peer and traffic-engineered link information.</p> <p>peer <name <i>name</i>>—(Optional) Display information for all peers or for the specified peer only.</p> <p>resource <name <i>name</i>>—(Optional) Display information for all resources or for the specified resource only.</p> <p>te-link <name <i>name</i>>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 213 • show link-management peer on page 217 • show link-management statistics on page 222 • show link-management te-link on page 224
List of Sample Output	show link-management routing on page 221
Output Fields	Table 22 on page 219 describes the output fields for the show link-management routing command. Output fields are listed in the approximate order in which they appear.

Table 22: show link-management routing Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down.
Control address	Address to which a control channel is established.
Control channel	Interface over which control packets are sent.

Table 22: show link-management routing Output Fields (*continued*)

Field Name	Field Description
State	State of the control channel.
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: Up or Down .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Resource	Forwarding adjacency LSP information.
Type	Type of resource. The type is always a forwarding adjacency LSP.
State	State of the LSP: Up or Down .
System Identifier	Internal identifier for the peer. The range of values is 0 through 64,000 .
Total bandwidth	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.
Traffic parameters	<ul style="list-style-type: none"> • Encoding—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET, Ethernet, and Packet. • Switching—Type of switching that can be performed on the traffic-engineered link: PSC-1 and Packet. • Granularity—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always unknown.

Sample Output

show link-management routing

```

user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
State: Up, Control address: (null)
Control-channel          State
fe-0/1/0.0               Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
State: Up, Control address: (null)
Control-channel          State
fe-0/1/2.0               Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
State: Down, Control address: (null)
Control-channel          State
so-0/2/0.0               State

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
State: Down, Control address: (null)
Control-channel          State
so-0/2/1.0               State

...

TE link name: __rpd:fe-0/1/0.0, State: Up
Local identifier: 2147483649, Remote identifier: 0,
Local address: 192.168.37.66, Remote address: 192.168.37.66,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:fe-0/1/2.0, State: Up
Local identifier: 2147483650, Remote identifier: 0,
Local address: 192.168.37.73, Remote address: 192.168.37.73,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:so-0/2/0.0, State: Down
Local identifier: 2147483651, Remote identifier: 0,
Local address: 192.168.37.82, Remote address: 192.168.37.95,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
Total bandwidth: 0bps, Traffic parameters: Encoding: Packet, Switching: Packet,
Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps, Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown

```

show link-management statistics

Syntax	show link-management statistics <peer <name <i>name</i> >>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display statistical information for Link Management Protocol (LMP) packets.
Options	none —Display information for all peers. peer <name <i>name</i>> —(Optional) Display information for all peers or for the specified peer only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 213 • show link-management peer on page 217 • show link-management routing on page 219 • show link-management te-link on page 224
List of Sample Output	show link-management statistics on page 223
Output Fields	Table 23 on page 222 describes the output fields for the show link-management statistics command. Output fields are listed in the approximate order in which they appear.

Table 23: show link-management statistics Output Fields

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.
Stale acknowledgments	Number of configAck and LinkSummaryAck packets received that have a stale message ID.

Table 23: show link-management statistics Output Fields (*continued*)

Field Name	Field Description
Stale negative acknowledgments	Number of configNack and LinkSummaryNack packets received that have a stale message ID.
Sent packets	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Retransmitted packets	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Dropped packets	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

Sample Output

show link-management statistics

```

user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
  Received packets
    Config: 1
    Hello: 2572
  Small packets: 0
  Wrong protocol version: 0
  Messages for unknown peer: 0
  Messages for bad state: 0
  Stale acknowledgments: 0
  Stale negative acknowledgments: 0
  Sent packets
    Config: 2
    ConfigAck: 1
    Hello: 2572
  Retransmitted packets
    Config: 1

```

show link-management te-link

Syntax	show link-management te-link <brief detail> <name <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.
Options	none —Display information for all traffic-engineered links. brief detail —(Optional) Display the specified level of output. name <i>name</i> —(Optional) Display information for the specified traffic-engineered link only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 213 • show link-management peer on page 217 • show link-management routing on page 219 • show link-management statistics on page 222
List of Sample Output	show link-management te-link on page 225
Output Fields	Table 24 on page 224 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 24: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: Up or Down .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH .

Table 24: show link-management te-link Output Fields (*continued*)

Field Name	Field Description
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
Available Bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Name	Name of the interface.
State	State of the interface: Up or Down .
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: Yes or No .
LSP-name	LSP name.

Sample Output

show link-management te-link

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-bd   Dn      43077      0           0bps No
TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
  Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
  Available bandwidth: 8Mbps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-be   Up      43076      0          10Mbps Yes  e2elasp-bf

```

show mpls admin-groups

Syntax	show mpls admin-groups <logical-system (all <i>logical-system-name</i>)>
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.
Description	Display information about configured Multiprotocol Label Switching (MPLS) administrative groups.
Options	none —Display information about the configured MPLS administrative groups. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show mpls admin-groups on page 226
Output Fields	Table 25 on page 226 describes the output fields for the show mpls admin-groups command. Output fields are listed in the approximate order in which they appear.

Table 25: show mpls admin-groups Output Fields

Field Name	Field Description
Group	Name of the administrative group.
Bit index	Value assigned to the administrative group.

Sample Output

show mpls admin-groups

```

user@host> show mpls admin-groups
Group      Bit index
black      3
blue       2
gold       1
green      0

```

show mpls call-admission-control

Syntax	show mpls call-admission-control <logical-system (all <i>logical-system-name</i>)> < <i>lsp-name</i> >
Syntax (EX Series Switches)	show mpls call-admission-control < <i>lsp-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.
Options	<p>none—Display CAC information for all LSPs.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>lsp-name</i>—(Optional) Display CAC information for the specified LSP only.</p>
Additional Information	The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.
Required Privilege Level	view
List of Sample Output	show mpls call-admission-control on page 228
Output Fields	Table 26 on page 227 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 26: 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.

Table 26: show mpls call-admission-control Output Fields (*continued*)

Field Name	Field Description
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

Sample Output

show mpls call-admission-control

```

user@host# show mpls call-admission-control

LSP name: pro1-be
*Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
*Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
*Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
*Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
*Primary  prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby  sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>

```

show mpls cspf

Syntax	show mpls cspf <logical-system (all <i>logical-system-name</i>)>
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.
Description	Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.
Options	none —Display MPLS CSFP statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show mpls cspf on page 230
Output Fields	Table 27 on page 229 describes the output fields for the show mpls cspf command. Output fields are listed in the approximate order in which they appear.

Table 27: 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.

Table 27: show mpls cspf Output Fields (*continued*)

Field Name	Field Description
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
Paths          total      successful    no route    sys error    CSPFs
               0          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

Syntax	show mpls diffserve-te <logical-system (all <i>logical-system-name</i>)>
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.
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. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show mpls diffserv-te on page 232
Output Fields	Table 28 on page 231 describes the output fields for the show mpls diffserv-te command. Output fields are listed in the approximate order in which they appear.

Table 28: show mpls diffserv-te Output Fields

Field Name	Field Description
Bandwidth model	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
TE class	DiffServ traffic engineering class.
Traffic class	MPLS class type that corresponds to the DiffServ traffic engineering class: <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.

Sample Output

`show mpls diffserv-te`

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class    Traffic class    Priority
te0         ct0              3
te1         ct1              2
```

show route forwarding-table

Syntax	<pre>show route forwarding-table <detail extensive summary> <ccc ccc-interface-name> <destination> <family family-name> <label label> <matching ip_prefix> <multicast> <vpn vpn></pre>
Release Information	Command introduced in Junos OS Release 9.5 for EX Series switches.
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	<p>none—Display the routes in the forwarding table.</p> <p>detail extensive summary—(Optional) Display the specified level of output.</p> <p>ccc—(Optional) Display the specified circuit cross-connect interface name for entries to match.</p> <p>destination —(Optional) Display the destination prefix.</p> <p>family family-name —(Optional) Display routing table entries for the specified family: ethernet-switching, inet, inet6, iso, mpls, vlan classification.</p> <p>label label —(Optional) Display route entries for the specified label name.</p> <p>matching ip_prefix —(Optional) Display route entries for the specified IP prefix.</p> <p>multicast—(Optional) Display route entries for multicast routes.</p> <p>vpn vpn —(Optional) Display route entries for the specified VPN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring MPLS on EX Series Switches on page 29 • Configuring MPLS on Provider Edge Switches (CLI Procedure) • Configuring MPLS on Provider Switches (CLI Procedure) on page 80
List of Sample Output	<p>show route forwarding-table on page 235</p> <p>show route forwarding-table summary on page 236</p> <p>show route forwarding-table extensive on page 236</p>

[show route forwarding-table ccc on page 238](#)
[show route forwarding-table family on page 238](#)
[show route forwarding-table label on page 239](#)
[show route forwarding-table matching on page 239](#)
[show route forwarding-table multicast on page 239](#)

Output Fields Table 29 on page 234 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 29: 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
Route Type (Type)	How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> cloned (clon)—(TCP or multicast only) Cloned route. destination (dest)—Remote addresses directly reachable through an interface. destination down (iddn)—Destination route for which the interface is unreachable. interface cloned (ifcl)—Cloned route for which the interface is unreachable. route down (ifdn)—Interface route for which the interface is unreachable. ignore (ignr)—Ignore this route. interface (intf)—Installed as a result of configuring an interface. permanent (perm)—Routes installed by the kernel when the routing table is initialized. user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route reference (RtRef)	Number of routes to reference.	detail , extensive
Flags	Route type flags: <ul style="list-style-type: none"> none—No flags are enabled. accounting—Route has accounting enabled. cached—Cache route. incoming-iface interface-number—Check against incoming interface. prefix load balance—Load balancing is enabled for this prefix. 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 29: 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 (mcrst)—Regular multicast next hop • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp) —Multicast group member. • receive (rcv)—Receive. • reject (rjct) Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail, extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail, extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	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
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	loc1	1308	2	
2.2.2.2/32	dest	0	2.2.2.2	loc1	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	loc1	1312	2	
3.3.3.1/32	dest	0	3.3.3.1	loc1	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	loc1	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	loc1	322	2	
10.93.13.238/32	dest	0	10.93.13.238	loc1	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	loc1	1318	2	
14.14.14.2/32	iddn	0	14.14.14.2	loc1	1318	2	
14.14.14.255/32	iddn	0	14.14.14.255	bcst	1316	1	ge-0/0/25.0
224.0.0.0/4	perm	1		mdsc	35	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	31	3	
224.0.0.5/32	user	1	224.0.0.5	mcst	31	3	
255.255.255.255/32	perm	0		bcst	32	1	

show route forwarding-table summary

```
user@switch> show route forwarding-table summary
```

```
Routing table: default.inet
```

```
Internet:
```

```

user:          6 routes
perm:          5 routes
intf:          8 routes
dest:         12 routes
ifdn:          1 routes
iddn:          3 routes
```

show route forwarding-table extensive

```
user@switch> show route forwarding-table summary
```

```
Routing table: default.inet [Index 0]
```

```
Internet:
```

```
Destination: default
```

```
Route type: user
```

```
Route reference: 2
```

```
Route interface-index: 0
```

```

Flags: sent to PFE, rt nh decoupled
Nexthop: 0:12:f2:21:cf:0
Next-hop type: unicast          Index: 333      Reference: 5
Next-hop interface: me0.0

Destination: default
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: none
Next-hop type: reject          Index: 36      Reference: 2

Destination: 0.0.0.0/32
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Next-hop type: discard         Index: 34      Reference: 1

Destination: 2.2.2.0/24
Route type: interface
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Next-hop type: resolve         Index: 1309    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.0/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.0
Next-hop type: receive         Index: 1307    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.1/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 0:21:59:cc:89:c0
Next-hop type: unicast         Index: 1320    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.2/32
Route type: interface
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308    Reference: 2

Destination: 2.2.2.2/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: none
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308    Reference: 2

Destination: 2.2.2.255/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.255

```

Next-hop type: broadcast Index: 1306 Reference: 1
 Next-hop interface: ae0.0

show route forwarding-table ccc

```
user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
ge-0/0/0.10      (CCC) user    0 3.3.3.2      Push 300112 1343    2 ae1.0
```

show route forwarding-table family

```
user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm    0
0                user    0      recv    49    3
1                user    0      recv    49    3
2                user    0      recv    49    3
299776           user    0      Pop     1334   2 ge-0/0/0.10
299792           user    0      Pop     1339   2 ge-0/0/0.14
299808           user    0      Pop     1341   2 ge-0/0/0.2
299824           user    0      Pop     1344   2 ge-0/0/0.11
299840           user    0      Pop     1345   2 ge-0/0/0.13
299856           user    0      Pop     1346   2 ge-0/0/0.18
299872           user    0      Pop     1347   2 ge-0/0/0.16
299888           user    0      Pop     1348   2 ge-0/0/0.7
299904           user    0      Pop     1349   2 ge-0/0/0.20
299920           user    0      Pop     1350   2 ge-0/0/0.19
299936           user    0      Pop     1351   2 ge-0/0/0.17
299952           user    0      Pop     1352   2 ge-0/0/0.9
299968           user    0      Pop     1353   2 ge-0/0/0.1
299984           user    0      Pop     1354   2 ge-0/0/0.12
300000           user    0      Pop     1355   2 ge-0/0/0.8
300016           user    0      Pop     1356   2 ge-0/0/0.4
300032           user    0      Pop     1357   2 ge-0/0/0.5
300048           user    0      Pop     1358   2 ge-0/0/0.3
300064           user    0      Pop     1359   2 ge-0/0/0.15
ge-0/0/0.1       (CCC) user    0 3.3.3.2      Push 300064 1340    2 ae1.0
ge-0/0/0.2       (CCC) user    0 3.3.3.2      Push 299872 1328    2 ae1.0
ge-0/0/0.3       (CCC) user    0 3.3.3.2      Push 299792 1323    2 ae1.0
ge-0/0/0.4       (CCC) user    0 3.3.3.2      Push 300016 1337    2 ae1.0
ge-0/0/0.5       (CCC) user    0 3.3.3.2      Push 299824 1325    2 ae1.0
ge-0/0/0.7       (CCC) user    0 3.3.3.2      Push 299920 1331    2 ae1.0
ge-0/0/0.8       (CCC) user    0 3.3.3.2      Push 299840 1326    2 ae1.0
ge-0/0/0.9       (CCC) user    0 3.3.3.2      Push 299888 1329    2 ae1.0
ge-0/0/0.10      (CCC) user    0 3.3.3.2      Push 300112 1343    2 ae1.0
ge-0/0/0.11      (CCC) user    0 3.3.3.2      Push 299776 1322    2 ae1.0
ge-0/0/0.12      (CCC) user    0 3.3.3.2      Push 299952 1333    2 ae1.0
ge-0/0/0.13      (CCC) user    0 3.3.3.2      Push 300096 1342    2 ae1.0
ge-0/0/0.14      (CCC) user    0 3.3.3.2      Push 299984 1335    2 ae1.0
ge-0/0/0.15      (CCC) user    0 3.3.3.2      Push 299936 1332    2 ae1.0
ge-0/0/0.16      (CCC) user    0 3.3.3.2      Push 299808 1324    2 ae1.0
ge-0/0/0.17      (CCC) user    0 3.3.3.2      Push 300000 1336    2 ae1.0
ge-0/0/0.18      (CCC) user    0 3.3.3.2      Push 300032 1338    2 ae1.0
ge-0/0/0.19      (CCC) user    0 3.3.3.2      Push 299904 1330    2 ae1.0
ge-0/0/0.20      (CCC) user    0 3.3.3.2      Push 299856 1327    2 ae1.0
```


show route forwarding-table label

```
user@switch> show route forwarding-table label 29976
```

```
Routing table: default.mpls
```

```
MPLS:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
29976	user	0		Pop	1334	2	ge-0/0/0.10

show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3
```

```
Routing table: default.inet
```

```
Internet:
```

show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast
```

```
Routing table: default.inet
```

```
Internet:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
224.0.0.0/4	perm	1		mdsc	35	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	31	3	
224.0.0.5/32	user	1	224.0.0.5	mcst	31	3	

```
Routing table: __master.anon__.inet
```

```
Internet:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
224.0.0.0/4	perm	0		mdsc	1289	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	1285	1	

```
Routing table: default.inet6
```

```
Internet6:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

show mpls interface

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 <i>interface-name</i> and set interfaces <i>interface-name</i> unit 0 family mpls commands.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Example: Configuring MPLS on EX Series Switches on page 29• Configuring MPLS on Provider Edge Switches (CLI Procedure)• Configuring MPLS on Provider Switches (CLI Procedure) on page 80
List of Sample Output	show mpls interface on page 240
Output Fields	Table 30 on page 240 describes the output fields for the show mpls interface command. Output fields are listed in the approximate order in which they appear.

Table 30: 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 interface

Syntax	show mpls interface <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show mpls interface
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.
Options	none —Display information about MPLS-enabled interfaces. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Additional Information	MPLS is enabled on an interface when the interface is configured with both the set protocol mpls interface <i>interface-name</i> and set interface <i>interface-name</i> unit 0 family mpls statements.
Required Privilege Level	view
List of Sample Output	show mpls interface on page 242
Output Fields	Table 31 on page 241 describes the output fields for the show mpls interface command. Output fields are listed in the approximate order in which they appear.

Table 31: 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.

Table 31: show mpls interface Output Fields (*continued*)

Field Name	Field Description
Always mark connection protection tlv	Enabled or Disabled: Enabled indicates that the <code>always-mark-connection-protection-tlv</code> statement is configured at the <code>[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static]</code> or the <code>[edit protocols mpls interface <i>interface-name</i> static]</code> 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 <code>switch-away-lsps</code> statement must be configured.
Switch away lsps	Enabled or Disabled: Enabled indicates that the <code>switch-away-lsps</code> statement is configured at the <code>[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static]</code> or the <code>[edit protocols mpls interface <i>interface-name</i> static]</code> hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

Sample Output

show mpls interface

```

user@host> show mpls interface

Interface: ge-0/2/1.57
State: Up
Administrative group: <none>
Maximum labels: 5
Static protection revert time: 5 seconds
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled

```

show mpls lsp

Syntax	<pre>show mpls lsp <brief detail extensive terse> <bidirectional unidirectional> <bypass> <count-active-routes> <defaults> <descriptions> <down up> <logical-system (all <i>logical-system-name</i>)> <lsp-type> <name <i>name</i>> <p2mp> <statistics> <transit></pre>
Syntax (EX Series Switches)	<pre>show mpls lsp <brief detail extensive terse> <bidirectional unidirectional> <bypass> <descriptions> <down up> <lsp-type> <name <i>name</i>> <p2mp> <statistics> <transit></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>defaults option added in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p>
Description	Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).
Options	<p>none—Display standard information about all configured and active dynamic MPLS LSPs.</p> <p>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.</p> <p>bidirectional unidirectional—(Optional) Display bidirectional or unidirectional LSP information, respectively.</p> <p>bypass—(Optional) Display LSPs used for protecting other LSPs.</p> <p>count-active-routes—(Optional) Display active routes for LSPs.</p> <p>defaults—(Optional) Display the MPLS LSP default settings.</p> <p>descriptions—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the [edit</p>

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.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified LSP or group of LSPs.

p2mp—(Optional) Display information about point-to-multipoint LSPs.

statistics—(Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.



.....
NOTE: If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored.

When used with the **bypass** option (**show mpls lsp bypass statistics**), display statistics for the traffic that flows only through the bypass LSP.
.....

transit—(Optional) Display LSPs transiting this routing device.

Required Privilege Level

view

Related Documentation

- [clear mpls lsp on page 184](#)

List of Sample Output

[show mpls lsp defaults on page 251](#)
[show mpls lsp descriptions on page 251](#)
[show mpls lsp detail on page 251](#)
[show mpls lsp extensive on page 252](#)
[show mpls lsp ingress extensive on page 253](#)

[show mpls lsp p2mp on page 253](#)

[show mpls lsp p2mp detail on page 254](#)

[show mpls lsp detail count-active-routes on page 254](#)

[show mpls lsp statistics extensive on page 255](#)

Output Fields [Table 32 on page 245](#) describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

Table 32: 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:vpls:router-id:routing-instance-name</i> . The <i>identifier</i> is automatically generated by Junos OS.	All levels
P2MP branch count	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
P	An asterisk (*) under this heading indicates that the LSP is a primary path.	All levels
address	(detail and extensive) Destination (egress routing device) of the LSP.	detail extensive
To	Destination (egress routing device) of the session.	brief
From	Source (ingress routing device) of the session.	brief detail
State	State of the LSP handled by this RSVP session: Up , Dn (down), or Restart .	brief detail
Active Route	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail extensive
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	brief
ActivePath	(Ingress LSP) Name of the active path: Primary or Secondary .	detail extensive
LSPname	Name of the LSP.	brief detail

Table 32: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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
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

Table 32: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 node suggests to use in the Resv message that is sent.	All levels
Suggested label sent	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
Autobandwidth	(Ingress LSP) The LSP is performing autobandwidth allocation.	detail extensive
MinBW	(Ingress LSP) Configured minimum value of the LSP, in bps.	detail extensive
MaxBW	(Ingress LSP) Configured maximum value of the LSP, in bps.	detail extensive
AdjustTimer	(Ingress LSP) Configured value of the bandwidth adjustment timer, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	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
Bandwidth Adjustment in <i>nnn</i> second(s)	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	detail extensive
Underflow limit	(Ingress LSP) Configured value of the threshold underflow limit.	detail extensive
Underflow sample count	(Ingress LSP) Current value for the underflow sample count.	detail extensive
Underflow Max AvgBW	(Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.	detail extensive

Table 32: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active path indicator	(Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path. *Primary long Standby short	detail extensive
Primary	(Ingress LSP) Name of the primary path.	detail extensive
Secondary	(Ingress LSP) Name of the secondary path.	detail extensive
Standby	(Ingress LSP) Name of the path in standby mode.	detail extensive
State	(Ingress LSP) State of the path: Up or Dn (down).	detail extensive
COS	(Ingress LSP) Class-of-service value.	detail extensive
Bandwidth per class	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	detail extensive
Priorities	(Ingress LSP) Configured value of the setup priority and the hold priority respectively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority).	detail extensive
OptimizeTimer	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
SmartOptimizeTimer	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
Reoptimization in xxx seconds	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	detail extensive
Computed ERO (S [L] denotes strict [loose] hops)	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	detail extensive
CSPF metric	(Ingress LSP) Constrained Shortest Path First metric for this path.	detail extensive

Table 32: 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. • 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08. 	detail extensive
Index number	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	extensive
Date	(Ingress LSP) Date of the LSP event.	extensive
Time	(Ingress LSP) Time of the LSP event.	extensive
Event	(Ingress LSP) Description of the LSP event.	extensive
Created	(Ingress LSP) Date and time the LSP was created.	extensive
Resv style	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail extensive
Labelin	Incoming label for this LSP.	brief detail
Labelout	Outgoing label for this LSP.	brief detail
LSPname	Name of the LSP.	brief detail

Table 32: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Time left	Number of seconds remaining in the lifetime of the reservation.	detail
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
Port number	Protocol ID and sender or receiver port used in this RSVP session.	detail
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	detail
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	detail
RESV rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the PATH rcvfrom field, indicates that the RSVP negotiation is complete.	detail
Record route	Recorded route for the session, taken from the record route object.	detail
Soft preempt	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	detail
Soft preemption pending	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	detail
MPLS-TE LSP Defaults	Default settings for MPLS traffic engineered LSPs: <ul style="list-style-type: none"> • LSP Holding Priority—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. • LSP Setup Priority—Determines whether a new LSP that preempts an existing LSP can be established. • Hop Limit—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress). • Bandwidth—Specifies the bandwidth in bits per second for the LSP. • LSP Retry Timer—Length of time in seconds that the ingress router waits between attempts to establish the primary path. 	defaults

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth**. You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.

Sample Output

show mpls lsp defaults

```
user@host> show mpls lsp defaults
MPLS-TE LSP Defaults
  LSP Holding Priority      0
  LSP Setup Priority       7
  Hop Limit                255
  Bandwidth                0
  LSP Retry Timer          30 seconds
```

show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To          LSP name          Description
10.0.0.195  to-sanjose                to-sanjose-desc
10.0.0.195  to-sanjose-other-desc      other-desc
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
  10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
      10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

show mpls lsp extensive

user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
ActivePath: (primary)
LSPtype: Static Configured, Ultimate hop popping
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Priorities: 7 0
SmartOptimizeTimer: 180
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt 20=Node-ID):
10.0.0.18 10.0.0.22
11 Sep 20 15:54:35.032 Make-before-break: Switched to new instance
10 Sep 20 15:54:34.029 Record Route: 10.0.0.18 10.0.0.22
9 Sep 20 15:54:34.029 Up
8 Sep 20 15:54:20.271 Originate make-before-break call
7 Sep 20 15:54:20.271 CSPF: computation result accepted 10.0.0.18 10.0.0.22

6 Sep 20 15:52:10.247 Selected as active path
5 Sep 20 15:52:10.246 Record Route: 10.0.0.18 10.0.0.22
4 Sep 20 15:52:10.243 Up
3 Sep 20 15:52:09.745 Originate Call
2 Sep 20 15:52:09.745 CSPF: computation result accepted 10.0.0.18 10.0.0.22

1 Sep 20 15:51:39.903 CSPF failed: no route toward 192.168.0.4
Created: Thu Sep 20 15:51:08 2012
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
LSPname: E-D, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 148, Since: Thu Sep 20 15:52:10 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 49601 protocol 0
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 27 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

show mpls lsp ingress extensive

```

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

50.0.0.1
  From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPname: test
  ActivePath: (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    OptimizeTimer: 300
    SmartOptimizeTimer: 180
    Reoptimization in 240 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
    1.1.1.2 S 4.4.4.1 S 5.5.5.2 S
      Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
      20=Node-ID):
        1.1.1.2 4.4.4.1 5.5.5.2
      17 Aug 3 13:17:33.601 CSPF: computation result ignored, new path less avail
      bw[3 times]
      16 Aug 3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2
      times]
      15 Aug 3 12:54:36.678 Selected as active path
      14 Aug 3 12:54:36.676 Record Route: 1.1.1.2 4.4.4.1 5.5.5.2
      13 Aug 3 12:54:36.676 Up
      12 Aug 3 12:54:33.924 Deselected as active
      11 Aug 3 12:54:33.924 Originate Call
      10 Aug 3 12:54:33.923 Clear Call
      9 Aug 3 12:54:33.923 CSPF: computation result accepted 1.1.1.2 4.4.4.1
      5.5.5.2
      8 Aug 3 12:54:33.922 2.2.2.2: No Route toward dest
      7 Aug 3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4
      times]
      6 Aug 3 12:35:03.830 Selected as active path
      5 Aug 3 12:35:03.828 Record Route: 2.2.2.2 3.3.3.2
      4 Aug 3 12:35:03.827 Up
      3 Aug 3 12:35:03.814 Originate Call
      2 Aug 3 12:35:03.814 CSPF: computation result accepted 2.2.2.2 3.3.3.2
      1 Aug 3 12:34:34.921 CSPF failed: no route toward 50.0.0.1
    Created: Tue Aug 3 12:34:35 2010
  Total 1 displayed, Up 1, Down 0

```

show mpls lsp p2mp

```

user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1          p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1          p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp p2mp detail

```
user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp1
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
  192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp2
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
  192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

      192.168.208.17
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail count-active-routes

```
user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

213.119.192.2
  From: 156.154.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
  ActivePath: (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 5Mbps MaxBW: 250Mbps
  AdjustTimer: 300 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 102 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Priorities: 7 0
    Bandwidth: 5Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
  10.252.0.177 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
    20=Node-ID):
```



```

10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp statistics extensive

```

user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPName: E-D
  Statistics: Packets 302, Bytes 28992
  Aggregate statistics: Packets 302, Bytes 28992
  ActivePath: (primary)
  LSPType: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
      6 Oct  3 11:18:28.281 Selected as active path
      5 Oct  3 11:18:28.281 Record Route:  10.0.0.18 10.0.0.22
      4 Oct  3 11:18:28.280 Up
      3 Oct  3 11:18:27.995 Originate Call
      2 Oct  3 11:18:27.995 CSPF: computation result accepted  10.0.0.18 10.0.0.22

      1 Oct  3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
  Created: Wed Oct  3 11:17:01 2012
Total 1 displayed, Up 1, Down 0

```

show mpls path

Syntax	show mpls path <logical-system (all <i>logical-system-name</i>)> < <i>path-name</i> >
Syntax (EX Series Switches)	show mpls path < <i>path-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).
Options	<p>none—Display standard information about all MPLS LSPs.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>path-name</i>—(Optional) Display information about the specified LSP only.</p>
Required Privilege Level	view
List of Sample Output	show mpls path on page 256
Output Fields	Table 33 on page 256 describes the output fields for the show mpls path command. Output fields are listed in the approximate order in which they appear.

Table 33: 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 rsvp interface

Syntax	show rsvp interface <brief detail extensive> <link-management> <logical-system (all <i>logical-system-name</i>)>
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.
Description	Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.
Options	<p>none—Display standard information about the status of RSVP-enabled interfaces and packet statistics.</p> <p>brief detail extensive link-management—(Optional) Display the specified level of output.</p> <p>link-management—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show rsvp interface brief on page 260 show rsvp interface detail on page 260 show rsvp interface extensive on page 260 show rsvp interface link-management on page 261
Output Fields	Table 34 on page 257 lists the output fields for the show rsvp interface command. Output fields are listed in the approximate order in which they appear.

Table 34: 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

Table 34: show RSVP interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	State of the interface. <ul style="list-style-type: none"> • Disabled—No traffic engineering information is displayed. • Down—Interface is not operational. • Enabled—Displays traffic engineering information. • Up—Interface is operational. 	All levels
NoAuthentication	Interface does not support RSVP authentication.	detail
NoAggregate	Interface does not support refresh reduction.	detail
NoReliable	Interface does not support refresh reduction message ID extension.	detail
NoLinkProtection	Interface does not support link protection.	detail
HelloInterval	Frequency at which RSVP hellos are sent on this interface (in seconds).	detail
Address	IP address of the local interface.	detail
Active control channel	Next-hop link address to transmit messages.	None specified
TElink	Traffic-engineered links that are managed by the peer they are associated with.	None specified
Active resv	Number of reservations that are actively reserving bandwidth on the interface.	All levels
PreemptionCnt	Number of times an RSVP session was preempted on this interface.	detail
Update threshold	Percentage change in reserved bandwidth to trigger an IGP update.	detail
Subscription	User-configured subscription factor.	All levels
bc number	Bandwidth allocated for the specified bandwidth constraint.	extensive
ct number	Bandwidth allocated for the specified class type.	extensive
Static BW	Total interface bandwidth, in bps.	All levels
Available BW	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	all 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
Overbooked BW	Currently overbooked bandwidth, in bps, by class type (ct0 through ct3).	detail

Table 34: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Highwater mark	Highest bandwidth that has ever been reserved on this interface, in bps.	brief
PacketType	Type of RSVP packet.	detail
Total Sent	Total number of packets sent.	detail
Total Received	Total number of packets received since RSVP was enabled.	detail
Last 5 seconds Sent	Number of packets sent in the last 5 seconds.	detail
Last 5 seconds Received	Number of packets received in the last 5 seconds.	detail
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	detail
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	detail
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	detail
Resv	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	detail
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	detail
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.	detail
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	detail
Ack	Acknowledge message for refresh reductions.	detail
Srefresh	Summary refresh messages.	detail
EndtoEnd RSVP	Statistics for the number of end-to-end RSVP messages sent.	detail
Queue	CoS transmit queue number and its associated forwarding class designation.	extensive
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

Table 34: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>queue-priority-value</i>	Low, High, None, or Exact. None indicates no rate limiting. Exact indicates the queue transmits at the configured rate only.	extensive

Sample Output

show rsvp interface brief

```

user@host> show rsvp interface brief
RSVP interface: 1 active

```

Interface	State	Active resv	Subscription	Static BW	Available BW	Reserved BW	Highwater mark
de0.0	Up	1	23%	10Mbps	989.992kbps	1.31Mbps	1.31Mbps

show rsvp interface detail

```

user@host> show rsvp interface detail
so-0/1/1.0 Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  SoftPreemptionCnt1
  OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
  PacketType
    Total
    Last 5 seconds

```

	Sent	Received	Sent	Received
Path	16	0	1	0
PathErr	0	0	0	0
PathTear	1	0	0	0
Resv	0	11	0	1
ResvErr	0	0	0	0
ResvTear	0	0	0	0
Hello	66	67	1	1
Ack	0	0	0	0
Srefresh	0	0	0	0
EndtoEnd RSVP	0	0	0	0

...

show rsvp interface extensive

```

user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 192.168.213.22, 10.255.240.175
  ActiveResv 1, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
  bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps
  bc2 = (ct2+ct3), StaticBW 311.04Mbps
  bc3 = ct3, StaticBW 155.52Mbps
  ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
  ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

```

```

ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Queue      TxRate      Priority Exact
  0          155.52Mbps      25%    Low
  1          155.52Mbps      25%    Low
  2          155.52Mbps      25%    Low
  3          155.52Mbps      25%    Low

```

show rsvp interface link-management

```

user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C State: Up
Active Control Channel: so-0/1/0.0

TElink: TElnk1, Link ID: 37811
ActiveResv 0, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

TElink: TElnk2, Link ID: 37808
ActiveResv 1, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B State: Up
Active Control Channel: so-1/0/0.0

TElink: TElnkAB1, Link ID: 1598
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

TElink: TElnkAB2, Link ID: 1597
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

```

show rsvp neighbor

Syntax	show rsvp neighbor <brief detail> <logical-system (all <i>logical-system-name</i>)>
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.
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. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show rsvp neighbor on page 266 show rsvp neighbor detail on page 266
Output Fields	Table 35 on page 262 lists the output fields for the show rsvp neighbor command. Output fields are listed in the approximate order in which they appear.

Table 35: 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.	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

Table 35: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Up cnt and Down cnt	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	detail
status	State of the RSVP neighbor: <ul style="list-style-type: none"> • Up—Routing device can detect RSVP Hello messages from the neighbor. • Down—Routing device has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP Hello messages sent by the neighbor. • Restarting—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled. • Restarted—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures. • Dead—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down. 	detail
LastChange	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> .	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 35: 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. • no operational—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions. 	detail
Remote end	<p>Neighboring routing device's status with regard to refresh reduction:</p> <ul style="list-style-type: none"> • enabled—Remote routing device has requested refresh reduction during RSVP message exchanges. • disabled—Remote routing device does not require refresh reduction. 	detail
Ack-extension	<p>An RSVP refresh reduction extension:</p> <ul style="list-style-type: none"> • enabled—Both local and remote routing devices support the ack-extension (RFC 2961). • disabled—Remote routing device does not support the ack-extension. 	detail
Link protection	<p>Status of the MPLS fast reroute mechanism that protects traffic from link failure:</p> <ul style="list-style-type: none"> • enabled—Link protection feature has been turned on, protecting the neighbor with a bypass LSP. • disabled—No link protection feature has been enabled for this neighbor. 	detail
LSP name	Name of the bypass LSP.	detail
Bypass LSP	<p>Status of the bypass LSP. It can have the following values:</p> <ul style="list-style-type: none"> • does not exist—Bypass LSP is not available. • connecting—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment. • operational—Bypass LSP is up and running. • down—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path. 	detail
Backup routes	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	detail
Backup LSPs	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	detail
Bypass explicit route	Explicit route object's (ERO) path that is taken by the bypass LSP.	detail

Table 35: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Restart time	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	detail
Recovery time	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	detail

Sample Output

show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203   0 3/2    13:01      3   366/349
192.168.207.207   0 1/0    22:49      3   448/448
```

show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: ecstasy1 status: Up
  Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
  Message received: 632
  Hello: sent 673, received 656, interval 3 sec
  Remote instance: 0x6432838a, Local instance: 0x74b72e36
  Refresh reduction: operational
    Remote end: enabled, Ack-extension: enabled
  Link protection: enabled
    LSP name: Bypass_to_192.168.207.203
    Bypass LSP: operational, Backup routes: 1, Backup LSPs: 0
    Bypass explicit route: 192.168.207.207 192.168.207.224
  Restart time: 60000 msec, Recovery time: 0 msec
```

show rsvp session

Syntax	<pre>show rsvp session <brief detail extensive terse> <bidirectional unidirectional> <down up> <interface <i>interface-name</i>> <lsp-type> <name <i>session-name</i>> <session-type> <statistics> <te-link <i>te-link</i>></pre>
Release Information	Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display information about Resource Reservation Protocol (RSVP) sessions.
Options	<p>none—Display standard information about all RSVP sessions.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>bidirectional unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.</p> <p>down up—(Optional) Display only LSPs that are inactive or active, respectively.</p> <p>interface <i>interface-name</i>—(Optional) Display RSVP sessions for the specified interface only.</p> <p><i>lsp-type</i>—(Optional) Display information about RSVP sessions with regard to LSPs:</p> <ul style="list-style-type: none"> • bypass—Sessions used for bypass LSPs. • lsp—Sessions used to set up LSPs. • nolsp—Sessions not used to set up LSPs. <p>name <i>session-name</i>—(Optional) Display information about the named session.</p> <p><i>session-type</i>—(Optional) Display information about a particular session type:</p> <ul style="list-style-type: none"> • egress—Sessions that terminate on this switch. • ingress—Sessions that originate from this switch. • transit—Sessions that transit through this switch. <p>statistics—(Optional) Display packet statistics.</p> <p>te-link <i>te-link</i>—(Optional) Display sessions with reservations on the specified traffic-engineered link name.</p>
Required Privilege Level	view

- Related Documentation**
- [Example: Configuring MPLS on EX Series Switches on page 29](#)
 - [Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect \(CLI Procedure\) on page 89](#)
 - [Configuring MPLS on Provider Edge Switches Using IP Over MPLS \(CLI Procedure\) on page 85](#)
 - [Configuring MPLS on Provider Switches \(CLI Procedure\) on page 80](#)
 - [Junos OS MPLS Applications Configuration Guide](#)

List of Sample Output

- [show rsvp session on page 270](#)
- [show rsvp session statistics on page 270](#)
- [show rsvp session detail on page 270](#)
- [show rsvp session extensive on page 270](#)

Output Fields Table 36 on page 268 describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 36: 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
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

Table 36: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

show rsvp session 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 session

Syntax `show rsvp session`
 `<brief | detail | extensive | terse>`
 `<bidirectional | unidirectional>`
 `<bypass>`
 `<down | up>`
 `<interface interface-name>`
 `<logical-system (all | logical-system-name)>`
 `<lsp-type>`
 `<name session-name>`
 `<p2mp>`
 `<session-type>`
 `<statistics>`
 `<te-link te-link>`

Syntax (EX Series Switches) `show rsvp session`
 `<brief | detail | extensive | terse>`
 `<bidirectional | unidirectional>`
 `<bypass>`
 `<down | up>`
 `<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.

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.

bypass—(Optional) Display RSVP sessions for bypass LSPs.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

interface *interface-name*—(Optional) Display RSVP sessions for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.

- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

name session-name—(Optional) Display information about the named session.

p2mp—(Optional) Display point-to-multipoint information.

session-type—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

statistics—(Optional) Display packet statistics.

te-link te-link—(Optional) Display sessions with reservations on the specified TE link.

Required Privilege Level view

Related Documentation [• clear rsvp session on page 186](#)

List of Sample Output [show rsvp session on page 277](#)
[show rsvp session statistics on page 277](#)
[show rsvp session detail on page 277](#)
[show rsvp session detail \(Path MTU Output Field\) on page 278](#)
[show rsvp session detail \(GMPLS\) on page 278](#)
[show rsvp session extensive on page 279](#)
[show rsvp session p2mp \(Ingress Router\) on page 279](#)
[show rsvp session p2mp \(Transit Router\) on page 280](#)

Output Fields [Table 37 on page 273](#) describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 37: 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

Table 37: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
P2MP branch count	(Appears only when the p2mp option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels
To	Destination (egress routing device) of the session.	All levels
From	Source (ingress routing device) of the session.	All levels
State	State of the path: Up , Down , or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress routing device) of the LSP.	detail
From	Source (ingress routing device) of the session.	detail
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either Up , Dn (down), or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	brief detail
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
Active Route	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail
LSPname	Name of the LSP.	brief detail
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices. LSPpath can also indicate when a graceful LSP deletion has been triggered.	detail
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

Table 37: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	detail
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	detail
Resv style or Style	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail
Label in	Incoming label for this LSP.	brief detail
Label out	Outgoing label for this LSP.	brief detail
Time left	Number of seconds remaining in the lifetime of the reservation.	brief detail
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
DiffServ info	Indicates whether the LSP is a multiclass LSP (multiclass diffServ-TE LSP) or a Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP).	detail
bandwidth	Bandwidth for each class type (ct0 , ct1 , ct2 , or ct3).	detail
Port number	Protocol ID and sender/receiver port used in this RSVP session.	detail
Attrib flags	Non-PHP indicates that ultimate hop popping has been requested by the LSP using this RSVP session	extensive
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

Table 37: show RSVP session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Node/Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	detail extensive
Type	<p>LSP type:</p> <ul style="list-style-type: none"> • Link protected LSP—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Node/Link protected LSP—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Protection down—LSP is not currently protected. • Bypass LSP—LSP that is used to protect one or more user LSPs in case of link failure. • Backup LSP at Point-of-Local-Repair (PLR)—LSP that has been temporarily established to protect a user LSP at the ingress of a failed link. • Backup LSP at Merge Point (MP)—LSP that has been temporarily established to protect a user LSP at the egress of a failed link. 	detail extensive
New bypass	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	extensive
Link protection up, using <i>bypass-name</i>	Link protection (the bypass name is also displayed) has been activated for the LSP.	extensive
Creating backup LSP, link down	A link down event occurred, and traffic is being switched over to the bypass LSP.	extensive
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	extensive
Path mtu	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the allow-fragmentation statement configured at the [edit protocols mpls path-mtu] hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	detail
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

Table 37: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Explct route	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail
Record route	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail

Sample Output

show rsvp session

```

user@host> show rsvp session
Ingress RSVP: 1 sessions
To          From          State  Rt  Style  Labelin Labelout LSPname
10.255.245.214 10.255.245.212 AdminDn 0 1 FF      -    22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To          From          State  Rt  Style  Labelin Labelout LSPname
10.255.245.194 10.255.245.195 Up      0 1 FF    39811    -  Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up      0 1 FF      3        -  pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To          From          State  Rt  Style  Labelin Labelout LSPname
10.255.245.198 10.255.245.197 Up      0 1 SE    100000    3  pro3-de
Total 1 displayed, Up 1, Down 0

```

show rsvp session statistics

```

user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To          From          State  Packets  Bytes  LSPname
10.255.245.24 10.255.245.22 Up        0        0  pro3-bd
10.255.245.24 10.255.245.22 Up    44868  2333136  pro3-bd-2
Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To          From          State  Packets  Bytes  LSPname
10.255.245.22 10.255.245.24 Up        0        0  pro3-db
10.255.245.22 10.255.245.24 Up        0        0  pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show rsvp session detail

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3

```

```

Time left:    -, Since: Fri Mar 26 18:42:42 2004
Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
Port number: sender 1 receiver 15876 protocol 0
PATH rcvfrom: localclient
Adspec: sent MTU 1500
PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

show rsvp session detail (Path MTU Output Field)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPname: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
  Port number: sender 1 receiver 57843 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 4470
  Path mtu: received 4470, using 4458 for forwarding
  PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
  RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
  Explct route: 192.168.37.89
  Record route: <self> 192.168.37.89 192.168.37.87
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
    Detour adspec: sent MTU 1512
    Path mtu: received 1512, using 1500 for forwarding

```

show rsvp session detail (GMPLS)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
  Resv style: 0 -, Label in: -, Label out: -
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH MTU: received 0
  PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
  Explct route: 100.100.100.100 93.93.93.93
  Record route: <self> 100.100.100.100 93.93.93.93
  Total 1 displayed, Up 0, Down 1
  Egress RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
  Transit RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0

```


show rsvp session extensive

```

user@host> show rsvp session extensive
Ingress RSVP: 1 sessions

192.168.0.4
  From: 192.168.0.5, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299808
  Resv style: 1 FF, Label in: -, Label out: 299808
  Time left: -, Since: Thu Sep 20 15:54:20 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 61576 protocol 0
  Attrib flags: Non-PHP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.0.18 (lt-1/2/0.17) 41 pkts
  RESV rcvfrom: 10.0.0.18 (lt-1/2/0.17) 40 pkts
  Explct route: 10.0.0.18 10.0.0.22
  Record route: <self> 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 140, Since: Thu Sep 20 15:52:10 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 49601 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 44 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show rsvp session p2mp (Ingress Router)

```

user@host> show rsvp session p2mp
Ingress RSVP: 3 sessions
P2MP name: test, P2MP branch count: 1
  To          From          State   Rt Style Labelin Labelout LSPname
  10.255.10.95 10.255.10.2   Up      0  1 SE  -          3 to-pe1
P2MP name: test2, P2MP branch count: 2
  To          From          State   Rt Style Labelin Labelout LSPname
  10.255.10.23 10.255.10.2   Up      0  1 SE  -          299776 to-pe3
  10.255.10.16 10.255.10.2   Up      0  1 SE  -          299776 to-pe4
Total 3 displayed, Up 3, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

show rsvp session p2mp (Transit Router)

user@host> show rsvp session p2mp

Ingress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.23	10.255.10.95	Up	0	1 SE	-	299792	to-pe2

Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.95	10.255.10.2	Up	0	1 SE	3	-	to-pe1

Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions

P2MP name: test2, P2MP branch count: 2

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.23	10.255.10.2	Up	0	1 SE	299776	299808	to-pe3
10.255.10.16	10.255.10.2	Up	0	1 SE	299776	299856	to-pe4

Total 2 displayed, Up 2, Down 0

show rsvp statistics

Syntax	show rsvp statistics <logical-system (all <i>logical-system-name</i>)>
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.
Description	Display Resource Reservation Protocol (RSVP) packet and error statistics.
Options	none —Display RSVP packet and error statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear rsvp statistics on page 188
List of Sample Output	show rsvp statistics on page 283
Output Fields	Table 38 on page 281 describes the output fields for the show rsvp statistics command. Output fields are listed in the approximate order in which they appear.

Table 38: 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.

Table 38: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
Resv FF	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
Resv WF	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
Res SE	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.
ResvConf	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
Ack	Acknowledge message for refresh reductions.
SRefresh	Summary refresh messages.
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.
EndtoEnd RSVP	Statistics for the number of End-to-end RSVP messages.
Errors	Statistics about errored RSVP packets.
Rcv pkt bad length	The packet was not processed because its length is inappropriate.
Rcv pkt unknown type	The packet is not one of the well-known RSVP types, as defined in RFC 2205, <i>Resource ReSerVation Protocol (RSVP)</i> .
Rcv pkt bad version	The packet is not an RSVP version 1 packet.
Rcv pkt auth fail	The packet failed authentication checks.
Rcv pkt bad checksum	The RSVP checksum check failed.
Rcv pkt bad format	General packet processing failed because the packet was badly formed.
Memory allocation fail	An internal resource failure occurred.
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.

Table 38: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
PathErr to client	Number of PathErr packets delivered to the local client.
ResvErr to client	Number of ResvErr packets delivered to the local client.
Path timeout	Number of times the sender timed out because the path was removed.
Resv timeout	Number of times the receiver timed out because the reservation was removed.
Message out-of-order	Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.
Unknown ack msg	A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1.
Recv nack	If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again.
Recv duplicated msg-id	Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.
No TE-link to 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

```
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
Ack	682	1414	0	0
SRefresh	395198	236030	5	2
Hello	578809	578221	4	4
EndtoEnd RSVP	0	0	0	0
Errors	Total		Last 5 seconds	
Rcv pkt bad length	0		0	
Rcv pkt unknown type	0		0	
Rcv pkt bad version	0		0	
Rcv pkt auth fail	0		0	
Rcv pkt bad checksum	0		0	
Rcv pkt bad format	0		0	
Memory allocation fail	0		0	
No path information	10		0	
Resv style conflict	0		0	
Port conflict	0		0	
Resv no interface	0		0	
PathErr to client	38		0	
ResvErr to client	0		0	
Path timeout	8		0	
Resv timeout	57		0	
Message out-of-order	0		0	
Unknown ack msg	2978		0	
Recv nack	86		0	
Recv duplicated msg-id	5		0	
No TE-link to recv Hop	0		0	
Rcv pkt disabled interface	0		0	
Transmit buffer full	0		0	
Transmit failure	0		0	
Receive failure	0		0	
P2MP RESV discarded by appl	0		0	
Rate limit	306		0	
Err msg loop detected	0		0	

show rsvp version

Syntax	show rsvp version <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show rsvp version
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.
Options	none —Display RSVP protocol settings. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show rsvp version on page 286
Output Fields	Table 39 on page 285 describes the output fields for the show rsvp version command. Output fields are listed in the approximate order in which they appear.

Table 39: show rsvp version Output Fields

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: Enabled or Disabled .
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: Aggressive , Disabled , or Normal . The default is Normal .
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.
Graceful deleting timeout	Currently configured value for the graceful-deletion-timeout statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

Table 39: show rsvp version Output Fields (*continued*)

Field Name	Field Description
NSR Mode	Status of the nonstop active routing feature for RSVP on the restarting device: Disabled , Enabled/Master , or Enabled/Standby .
NSR State	<p>State of the nonstop active routing feature for RSVP on the restarting device.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • Idle • TE-link sync complete • Neighbor sync complete • Path state sync complete • Resv state sync complete • Bypass sync complete • Init sync complete
Setup protection	Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: Enabled or Disabled
Graceful restart	Status of the graceful restart feature for RSVP on the restarting routing device: Enabled or Disabled .
Restart helper mode	Status of the helper mode feature: Enabled or Disabled . When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
Maximum helper restart time	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.
Maximum helper recovery time	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
Restart time	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
Recovery time	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
P2p transit LSP nexthop mode	Point-to-point transit LSP nexthop mode on PTX Series devices. The possible values are Chained or Unchained
P2mp transit LSP nexthop mode	Point-to-multipoint transit LSP nexthop mode on PTX Series devices. The possible values are Chained or Unchained

Sample Output

show rsvp version

```
user@host> show rsvp version
```



```
Resource ReSerVation Protocol, version 1. rfc2205
  RSVP protocol:           Enabled
  R(refresh timer):        30 seconds
  K(keep multiplier):      3
  Preemption:              Normal
  Soft-preemption cleanup:  30 seconds
  Graceful deletion timeout: 30 seconds
  NSR mode:                Enabled/Master
  NSR state:               Init sync complete
  Setup protection:        Disabled
  Graceful restart:        Disabled
  Restart helper mode:     Enabled
  Maximum helper restart time: 20000 msec
  Maximum helper recovery time: 180000 msec
  Restart time:            0 msec
  P2p transit LSP nexthop mode: Unchained
  P2mp transit LSP nexthop mode: Unchained
```

show ted database

Syntax	show ted database <brief detail extensive> <logical-system (all <i>logical-system-name</i>)> < <i>system-name</i> >
Syntax (EX Series Switches)	show ted database <brief detail extensive> < <i>system-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.
Options	<p>none—Display standard information about all entries in the traffic engineering database.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>system-name</i>—(Optional) Display traffic engineering database information for a particular system.</p>
Required Privilege Level	view
List of Sample Output	show ted database brief on page 290 show ted database detail system-name on page 291 show ted database extensive on page 291
Output Fields	Table 40 on page 288 describes the output fields for the show ted database command. Output fields are listed in the approximate order in which they appear.

Table 40: 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

Table 40: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type	Type of node. It can be either Rtr (router) or Net (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: <ul style="list-style-type: none"> • IS-IS(1)—IS-IS Level 1. • IS-IS(2)—IS-IS Level 2. • OSPF (area-number)—OSPF from the specified area. 	All levels
To	Address on the far end of a link.	detail extensive
Local	Address of the local interface being used to reach the remote node.	detail extensive
Remote	Address of the interface on the remote node.	detail extensive
Metric	Configured traffic engineering 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 40: 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: 6 ISIS nodes 6 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
cheviot.00(123.456.1.10)         Rtr   383     1     1 IS-IS(2) IS-IS(1)
corriedale.00(123.456.1.11)      Rtr    36     2     0 IS-IS(2) IS-IS(1)
wolff.00(123.456.1.12)          Rtr   399     0     0 IS-IS(2) IS-IS(1)
perendale.00(123.456.1.13)       Rtr   385     2     0 IS-IS(2) IS-IS(1)

```

```
merino.00(123.456.1.14)    Rtr    379    1    3 IS-IS(2) IS-IS(1)
romney.00(123.456.1.15)   Rtr    427    0    2 IS-IS(2) IS-IS(1)
```

show ted database detail system-name

```
user@host> show ted database detail merino
TED database: 6 ISIS nodes 6 INET nodes
NodeID: merino.00(123.456.1.14)
  Type: Rtr, Age: 507 secs, LinkIn: 1, LinkOut: 3
  Protocol: IS-IS(2)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207

    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
  Protocol: IS-IS(1)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207

    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
```

show ted database extensive

```
user@host> show ted database extensive
TED database: 0 ISIS nodes 2 INET nodes
NodeID: 10.255.245.196
  Type: Rtr, Age: 46 secs, LinkIn: 1, LinkOut: 1
  Protocol: OSPF(0.0.0.0)
    To: 10.255.245.24, Local: 4.4.4.4, Remote: 5.5.5.5
    Metric: 1
    Static BW: 155.52Mbps
    Reservable BW: 155.52Mbps
    Available BW [TE-class] bps:
      [te0] 155.52Mbps   [te1] 155.52Mbps   [te2] 155.52Mbps   [te3] 155.52Mbps

      [te4] 155.52Mbps   [te5] 155.52Mbps   [te6] 155.52Mbps   [te7] 155.52Mbps

    Diffserv-TE BW model: Maximum allocation model
    Static BW [CT-class] bps:
      [ct0] 155.52Mbps   [ct1] 155.52Mbps   [ct2] 155.52Mbps   [ct3] 155.52Mbps

    Interface Switching Capability Descriptor(1):
      Switching type: PSC-1
      Encoding type: SDH/SONET
      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
      Minimum LSP BW: 155.52Mbps
      Interface MTU: 1285
    Interface Switching Capability Descriptor(2):
      Switching type: TDM
      Encoding type: SDH/SONET
      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
      Minimum LSP BW: 155.52Mbps
      Interface supports standard SONET/SDH
```

show ted link

Syntax	show ted link <brief detail> <logical-system (all <i>logical-system-name</i>)>
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.
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. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show ted link brief on page 293 show ted link detail on page 293
Output Fields	Table 41 on page 292 describes the output fields for the show ted link command. Output fields are listed in the approximate order in which they appear.

Table 41: 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

Table 41: show ted link Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels

Sample Output

show ted link brief

```

user@host> show ted link brief
TED link:
ID                               ->ID                               LocalPath LocalBW
cheviot.00(123.456.1.10)         merino.00(123.456.1.14)           0 0bps
merino.00(123.456.1.14)         corriedale.00(123.456.1.11)       0 0bps
merino.00(123.456.1.14)         perendale.00(123.456.1.13)       0 0bps
merino.00(123.456.1.14)         cheviot.00(123.456.1.10)         0 0bps
romney.00(123.456.1.15)         corriedale.00(123.456.1.11)       0 0bps
romney.00(123.456.1.15)         perendale.00(123.456.1.13)       0 0bps

```

show ted link detail

```

user@host> show ted link detail
TED link:
cheviot.00(123.456.1.10)->merino.00(123.456.1.14), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
merino.00(123.456.1.14)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
merino.00(123.456.1.14)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
merino.00(123.456.1.14)->cheviot.00(123.456.1.10), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
romney.00(123.456.1.15)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
romney.00(123.456.1.15)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
  localBW [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps

```

show ted protocol

Syntax	show ted protocol <brief detail> <logical-system (all <i>logical-system-name</i>)>
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.
Description	Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.
Options	<p>none—Display standard information about the protocols from which the traffic engineering database learned about its nodes.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show ted protocol on page 294
Output Fields	Table 42 on page 294 describes the output fields for the show ted protocol command. Output fields are listed in the approximate order in which they appear.

Table 42: 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)

