

MPLS for EX9200 Switches

Release

15.1



Modified: 2015-06-03

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

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

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

MPLS for EX9200 Switches

15.1

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

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xi
	Documentation and Release Notes	xi
	Supported Platforms	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Overview	
Chapter 1	Understanding LDP	3
	LDP Introduction	3
	Junos OS LDP Protocol Implementation	3
Chapter 2	Understanding LDP Operations	5
	LDP Operation	5
	Label Operations	5
Chapter 3	Understanding LDP Messaging	9
	LDP Message Types	9
	Discovery Messages	9
	Session Messages	10
	Advertisement Messages	10
	Notification Messages	10
Chapter 4	Understanding LDP Session Protection	11
	LDP Session Protection	11
Part 2	Configuring LDP	
Chapter 5	Minimum Configuration and LDP Basic Settings	15
	Minimum LDP Configuration	15
	Specifying the Transport Address Used by LDP	16
	Enabling and Disabling LDP	16
	Configuring the LDP Timer for Hello Messages	17
	Configuring the LDP Timer for Link Hello Messages	17
	Configuring the LDP Timer for Targeted Hello Messages	17

	Configuring the Delay Before LDP Neighbors Are Considered Down	18
	Configuring the LDP Hold Time for Link Hello Messages	18
	Configuring the LDP Hold Time for Targeted Hello Messages	19
	Configuring the Interval for LDP Keepalive Messages	19
	Configuring the LDP Keepalive Timeout	19
	Configuring LDP Route Preferences	20
	Enabling Strict Targeted Hello Messages for LDP	20
Chapter 6	Configuring the LDP Forwarding Equivalence Class (FEC)	21
	Configuring FEC Deaggregation	21
	Configuring Policers for LDP FECs	22
	Configuring LDP IPv4 FEC Filtering	23
Chapter 7	Configuring BFD for LDP	25
	Configuring BFD for LDP LSPs	25
	Configuring ECMP-Aware BFD for LDP LSPs	28
	Configuring a Failure Action for the BFD Session on an LDP LSP	28
	Configuring the Holddown Interval for the BFD Session	29
Chapter 8	Configuring Graceful Restart	31
	LDP Graceful Restart	31
	Configuring LDP Graceful Restart	32
	Enabling Graceful Restart	32
	Disabling LDP Graceful Restart or Helper Mode	32
	Configuring Reconnect Time	33
	Configuring Recovery Time and Maximum Recovery Time	34
Chapter 9	Configuring LDP Filtering and Apply Policies	35
	Filtering Inbound LDP Label Bindings	35
	Examples: Filtering Inbound LDP Label Bindings	36
	Filtering Outbound LDP Label Bindings	37
	Examples: Filtering Outbound LDP Label Bindings	38
	Configuring OAM Ingress Policies for LDP	39
	Firewall Filter Match Conditions for MPLS Traffic	39
	Configuring the Prefixes Advertised into LDP from the Routing Table	41
	Example: Configuring the Prefixes Advertised into LDP	41
Chapter 10	Configuring LDP Traceroute and Collecting Statistics	43
	Configuring LDP LSP Traceroute	43
	Collecting LDP Statistics	44
	LDP Statistics Output	44
	Disabling LDP Statistics on the Penultimate-Hop Router	45
	LDP Statistics Limitations	46
	Tracing LDP Protocol Traffic	46
	Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels	46
	Tracing LDP Protocol Traffic Within FECs	47
	Examples: Tracing LDP Protocol Traffic	48

Part 3	Configuration Statements and Operational Commands	
Chapter 11	Configuration Statements	53
	[edit protocols bgp] Hierarchy Level	54
	Common BGP Family Options	54
	Complete [edit protocols bgp] Hierarchy	55
	[edit protocols ldp] Hierarchy Level	61
	[edit protocols mpls] Hierarchy Level	63
	Complete [edit protocols mpls] Hierarchy	63
	allow-subnet-mismatch	64
	authentication-algorithm	65
	authentication-key (Protocols LDP)	67
	bfd-liveness-detection (Protocols LDP)	68
	deaggregate	69
	disable (Protocols LDP)	70
	dod-request-policy	71
	downstream-on-demand	71
	ecmp	72
	egress-policy	72
	explicit-null (Protocols LDP)	73
	export (Protocols LDP)	73
	failure-action (Protocols LDP)	74
	fec	75
	graceful-restart (Protocols LDP)	76
	hello-interval (Protocols LDP)	77
	helper-disable (LDP)	78
	holddown-interval	79
	hold-time (Protocols LDP)	80
	igp-synchronization	81
	import (Protocols LDP)	82
	ingress-policy	83
	interface (Protocols LDP)	84
	keepalive-interval	85
	keepalive-timeout	86
	l2-smart-policy	86
	label-withdrawal-delay	87
	ldp	88
	ldp-p2mp	91
	log-updown (Protocols LDP)	92
	make-before-break (LDP)	93
	maximum-neighbor-recovery-time	94
	no-forwarding	95
	oam (Protocols LDP)	96
	p2mp (Protocols LDP)	97
	periodic-traceroute	98
	policing (Protocols LDP)	100
	preference (Protocols LDP)	101
	reconnect-time	102
	recovery-time	103

	session (ldp)	104
	session-protection	105
	strict-targeted-hellos	105
	targeted-hello	106
	traceoptions (Protocols LDP)	107
	track-igp-metric	109
	traffic-statistics (Protocols LDP)	110
	transport-address	112
Chapter 12	Operational Commands	113
	ping mpls ldp	114
	show ldp database	117
	show ldp session	126
	show ldp traffic-statistics	132
	show ldp session	136

List of Figures

Part 1	Overview	
Chapter 2	Understanding LDP Operations	5
	Figure 1: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs	6
	Figure 2: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs	7

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 2	Configuring LDP	
Chapter 9	Configuring LDP Filtering and Apply Policies	35
	Table 3: from Operators That Apply to LDP Received-Label Filtering	35
	Table 4: to Operators for LDP Outbound-Label Filtering	37
	Table 5: Firewall Filter Match Conditions for MPLS Traffic	39
Part 3	Configuration Statements and Operational Commands	
Chapter 12	Operational Commands	113
	Table 6: show ldp database Output Fields	118
	Table 7: show ldp session Output Fields	126
	Table 8: show ldp traffic-statistics Output Fields	133
	Table 9: show ldp session Output Fields	136

About the Documentation

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

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 xiii defines notice icons used in this guide.

Table 1: Notice Icons

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

Table 2 on page xiii 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

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<pre>user@host> show chassis alarms</pre> <p>No alarms currently active</p>
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	<p>Configure the machine's domain name:</p> <pre>[edit] root@# set system domain-name domain-name</pre>
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 <code>[edit protocols ospf area area-id]</code> 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.	<pre>[edit] routing-options { static { route default { nexthop <i>address</i>; retain; } } }</pre>
;(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.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Understanding LDP on page 3](#)
- [Understanding LDP Operations on page 5](#)
- [Understanding LDP Messaging on page 9](#)
- [Understanding LDP Session Protection on page 11](#)

CHAPTER 1

Understanding LDP

- [LDP Introduction on page 3](#)
- [Junos OS LDP Protocol Implementation on page 3](#)

LDP Introduction

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

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

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

Junos OS LDP Protocol Implementation

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

CHAPTER 2

Understanding LDP Operations

- [LDP Operation on page 5](#)
- [Label Operations on page 5](#)

LDP Operation

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

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

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

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

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

Related Documentation

- [Logical Interfaces](#)

Label Operations

[Figure 1 on page 6](#) depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see *MPLS Label Overview*.) The shaded inner oval

represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

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

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

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





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

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



CHAPTER 3

Understanding LDP Messaging

- [LDP Message Types on page 9](#)
- [Discovery Messages on page 9](#)
- [Session Messages on page 10](#)
- [Advertisement Messages on page 10](#)
- [Notification Messages on page 10](#)

LDP Message Types

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

- [Discovery Messages on page 9](#)
- [Session Messages on page 10](#)
- [Advertisement Messages on page 10](#)
- [Notification Messages on page 10](#)

Discovery Messages

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

LDP uses the following discovery procedures:

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

whether to respond to or ignore the targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages to the initiating router.

Session Messages

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

Advertisement Messages

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

Notification Messages

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

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

CHAPTER 4

Understanding LDP Session Protection

- [LDP Session Protection on page 11](#)

LDP Session Protection

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

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

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

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

PART 2

Configuring LDP

- [Minimum Configuration and LDP Basic Settings on page 15](#)
- [Configuring the LDP Forwarding Equivalence Class \(FEC\) on page 21](#)
- [Configuring BFD for LDP on page 25](#)
- [Configuring Graceful Restart on page 31](#)
- [Configuring LDP Filtering and Apply Policies on page 35](#)
- [Configuring LDP Traceroute and Collecting Statistics on page 43](#)

CHAPTER 5

Minimum Configuration and LDP Basic Settings

- [Minimum LDP Configuration on page 15](#)
- [Specifying the Transport Address Used by LDP on page 16](#)
- [Enabling and Disabling LDP on page 16](#)
- [Configuring the LDP Timer for Hello Messages on page 17](#)
- [Configuring the Delay Before LDP Neighbors Are Considered Down on page 18](#)
- [Configuring the Interval for LDP Keepalive Messages on page 19](#)
- [Configuring the LDP Keepalive Timeout on page 19](#)
- [Configuring LDP Route Preferences on page 20](#)
- [Enabling Strict Targeted Hello Messages for LDP on page 20](#)

Minimum LDP Configuration

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

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

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

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

Specifying the Transport Address Used by LDP

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

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

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

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

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

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

Related Documentation

- [transport-address on page 112](#)

Enabling and Disabling LDP

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

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

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

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

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

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


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

Configuring the LDP Timer for Hello Messages

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

There are two types of LDP hello messages:

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

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 18](#).

Configuring the LDP Timer for Link Hello Messages

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

```
hello-interval seconds;
```

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

Configuring the LDP Timer for Targeted Hello Messages

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

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

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

Configuring the Delay Before LDP Neighbors Are Considered Down

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

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



NOTE: By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router might declare an LDP neighbor down that is still functioning normally. For more information, see [“Configuring the LDP Timer for Hello Messages” on page 17](#).

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an LDP peer router advertises a shorter hold time than the value you have configured, the peer router’s advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

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

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

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

Configuring the LDP Hold Time for Link Hello Messages

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

```
hold-time seconds;
```

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

Configuring the LDP Hold Time for Targeted Hello Messages

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

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

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

Configuring the Interval for LDP Keepalive Messages

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

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 18](#).

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

```
keepalive-interval seconds;
```

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

Configuring the LDP Keepalive Timeout

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

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

```
keepalive-timeout seconds;
```

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

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

Configuring LDP Route Preferences

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

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

```
preference preference;
```

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

Enabling Strict Targeted Hello Messages for LDP

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

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

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

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

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

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

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

CHAPTER 6

Configuring the LDP Forwarding Equivalence Class (FEC)

- [Configuring FEC Deaggregation on page 21](#)
- [Configuring Policers for LDP FECs on page 22](#)
- [Configuring LDP IPv4 FEC Filtering on page 23](#)

Configuring FEC Deaggregation

When an LDP egress router advertises multiple prefixes, the prefixes are bound to a single label and aggregated into a single forwarding equivalence class (FEC). By default, LDP maintains this aggregation as the advertisement traverses the network.

Normally, because an LSP is not split across multiple next hops and the prefixes are bound into a single LSP, load-balancing across equal-cost paths does not occur. You can, however, load-balance across equal-cost paths if you configure a load-balancing policy and deaggregate the FECs.

Deaggregating the FECs causes each prefix to be bound to a separate label and become a separate LSP.

To configure deaggregated FECs, include the **deaggregate** statement:

deaggregate;

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

For all LDP sessions, you can configure deaggregated FECs only globally.

Deaggregating a FEC allows the resulting multiple LSPs to be distributed across multiple equal-cost paths and distributes LSPs across the multiple next hops on the egress segments but installs only one next hop per LSP.

To aggregate FECs, include the **no-deaggregate** statement:

no-deaggregate;

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

For all LDP sessions, you can configure aggregated FECs only globally.

**Related
Documentation**

- *Configuring Load Balancing Across RSVP LSPs*
- *Configuring Protocol-Independent Load Balancing in Layer 3 VPNs*
- *Configuring VPLS Load Balancing*
- *Example: Load Balancing BGP Traffic*

Configuring Policers for LDP FECs

You can configure the Junos OS to track and police traffic for LDP FECs. LDP FEC policers can be used to do any of the following:

- Track or police the ingress traffic for an LDP FEC.
- Track or police the transit traffic for an LDP FEC.
- Track or police LDP FEC traffic originating from a specific forwarding class.
- Track or police LDP FEC traffic originating from a specific virtual routing and forwarding (VRF) site.
- Discard false traffic bound for a specific LDP FEC.

To police traffic for an LDP FEC, you must first configure a filter. Specifically, you need to configure either the **interface** statement or the **interface-set** statement at the **[edit firewall family protocol-family filter filter-name term term-name from]** hierarchy level. The **interface** statement allows you to match the filter to a single interface. The **interface-set** statement allows you to match the filter to multiple interfaces.

For more information on how to configure the **interface** statement, the **interface-set** statement, and policers for LDP FECs, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Once you have configured the filters, you need to include them in the **policing** statement configuration for LDP. To configure policers for LDP FECs, include the **policing** statement:

```
policing {  
  fec fec-address {  
    ingress-traffic filter-name;  
    transit-traffic filter-name;  
  }  
}
```

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

The **policing** statement includes the following options:

- **fec**—Specify the FEC address for the LDP FEC you want to police.
- **ingress-filter**—Specify the name of the ingress traffic filter.
- **transit-traffic**—Specify the name of the transit traffic filter.

Configuring LDP IPv4 FEC Filtering

By default, when a targeted LDP session is established, the Junos OS always exchanges both the IPv4 forwarding equivalence classes (FECs) and the Layer 2 circuit FECs over the targeted LDP session. For an LDP session to an indirectly connected neighbor, you might only want to export Layer 2 circuit FECs to the neighbor if the session was specifically configured to support Layer 2 circuits or VPLS.

In a mixed vendor network where all non-BGP prefixes are advertised into LDP, the LDP database can become large. For this type of environment, it can be useful to prevent the advertisement of IPv4 FECs over LDP sessions formed because of Layer 2 circuit or LDP VPLS configuration. Similarly, it can be useful to filter any IPv4 FECs received in this sort of environment.

If all the LDP neighbors associated with an LDP session are Layer 2 only, you can configure the Junos OS to advertise only Layer 2 circuit FECs by configuring the **l2-smart-policy** statement. This feature also automatically filters out the IPv4 FECs received on this session. If you have configured an explicit export or import policy, this feature is disabled.

If one of the LDP session's neighbors is formed because of a discovered adjacency or if the adjacency is formed because of an LDP tunneling configuration on one or more RSVP LSPs, the IPv4 FECs are advertised and received using the default behavior.

To prevent LDP from exporting IPv4 FECs over LDP sessions with Layer 2 neighbors only and to filter out IPv4 FECs received over such sessions, include the **l2-smart-policy** statement:

```
l2-smart-policy;
```

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

CHAPTER 7

Configuring BFD for LDP

- [Configuring BFD for LDP LSPs on page 25](#)
- [Configuring ECMP-Aware BFD for LDP LSPs on page 28](#)
- [Configuring a Failure Action for the BFD Session on an LDP LSP on page 28](#)
- [Configuring the Holddown Interval for the BFD Session on page 29](#)

Configuring BFD for LDP LSPs

You can configure Bidirectional Forwarding Detection (BFD) for LDP LSPs. The BFD protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the router stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of static routes, providing faster detection.

An error is logged whenever a BFD session for a path fails. The following shows how BFD for LDP LSP log messages might appear:

```
RPD_LDP_BFD_UP: LDP BFD session for FEC 10.255.16.14/32 is up
RPD_LDP_BFD_DOWN: LDP BFD session for FEC 10.255.16.14/32 is down
```

You can also configure BFD for RSVP LSPs, as described in *Configuring BFD for MPLS IPv4 LSPs*.

The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

To enable BFD for LDP LSPs, include the **oam** and **bfd-liveness-detection** statements:

```
oam {
  bfd-liveness-detection {
    detection-time threshold milliseconds;
```

```

ecmp;
failure-action {
    remove-nexthop;
    remove-route;
}
holddown-interval seconds;
ingress-policy ingress-policy-name;
minimum-interval milliseconds;
minimum-receive-interval milliseconds;
minimum-transmit-interval milliseconds;
multiplier detection-time-multiplier;
no-adaptation;
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
version (0 | 1 | automatic);
}
fec fec-address {
    bfd-liveness-detection {
        detection-time threshold milliseconds;
        ecmp;
        failure-action {
            remove-nexthop;
            remove-route;
        }
        holddown-interval milliseconds;
        ingress-policy ingress-policy-name;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    no-bfd-liveness-detection;
    periodic-traceroute {
        disable;
        exp exp-value;
        fanout fanout-value;
        frequency minutes;
        paths number-of-paths;
        retries retry-attempts;
        source address;
        ttl ttl-value;
        wait seconds;
    }
}
lsp-ping-interval seconds;
periodic-traceroute {
    disable;
    exp exp-value;

```

```

fanout fanout-value;
frequency minutes;
paths number-of-paths;
retries retry-attempts;
source address;
ttl ttl-value;
wait seconds;
}
}

```

You can enable BFD for the LDP LSPs associated with a specific forwarding equivalence class (FEC) by configuring the FEC address using the **fec** option at the **[edit protocols ldp]** hierarchy level. Alternatively, you can configure an Operation Administration and Management (OAM) ingress policy to enable BFD on a range of FEC addresses. For more information, see [“Configuring OAM Ingress Policies for LDP” on page 39](#).

You cannot enable BFD LDP LSPs unless their equivalent FEC addresses are explicitly configured or OAM is enabled on the FECs using an OAM ingress policy. If BFD is not enabled for any FEC addresses, the BFD session will not come up.

You can configure the **oam** statement at the following hierarchy levels:

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

The **oam** statement includes the following options:

- **fec**—Specify the FEC address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.
- **lsp-ping-interval**—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command. For more information, see the [CLI Explorer](#).

The **bfd-liveness-detection** statement includes the following options:

- **ecmp**—Cause LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the **ecmp** option, you must also configure the **periodic-traceroute** statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the **periodic-traceroute** statement at the global hierarchy level (**[edit protocols ldp oam]**) while only configuring the **ecmp** option for a specific FEC (**[edit protocols ldp oam fec address bfd-liveness-detection]**).
- **holddown-interval**—Specify the duration the BFD session should remain up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.
- **minimum-interval**—Specify the minimum transmit and receive interval. If you configure the **minimum-interval** option, you do not need to configure the **minimum-receive-interval** option or the **minimum-transmit-interval** option.
- **minimum-receive-interval**—Specify the minimum receive interval. The range is from 1 through 255,000 milliseconds.

- **minimum-transmit-interval**—Specify the minimum transmit interval. The range is from 1 through 255,000 milliseconds.
- **multiplier**—Specify the detection time multiplier. The range is from 1 through 255.
- **version**—Specify the BFD version. The options are BFD version 0 or BFD version 1. By default, the Junos OS software attempts to automatically determine the BFD version.

Configuring ECMP-Aware BFD for LDP LSPs

When you configure BFD for a FEC, a BFD session is established for only one active local next-hop for the router. However, you can configure multiple BFD sessions, one for each FEC associated with a specific equal-cost multipath (ECMP) path. For this to function properly, you also need to configure LDP LSP periodic traceroute. (See “[Configuring LDP LSP Traceroute](#)” on page 43.) LDP LSP traceroute is used to discover ECMP paths. A BFD session is initiated for each ECMP path discovered. Whenever a BFD session for one of the ECMP paths fails, an error is logged.

LDP LSP traceroute is run periodically to check the integrity of the ECMP paths. The following might occur when a problem is discovered:

- If the latest LDP LSP traceroute for a FEC differs from the previous traceroute, the BFD sessions associated with that FEC (the BFD sessions for address ranges that have changed from previous run) are brought down and new BFD sessions are initiated for the destination addresses in the altered ranges.
- If the LDP LSP traceroute returns an error (for example, a timeout), all the BFD sessions associated with that FEC are torn down.

To configure LDP to establish BFD sessions for all ECMP paths configured for the specified FEC, include the **ecmp** statement.

ecmp;

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

Along with the **ecmp** statement, you must also include the **periodic-traceroute** statement, either in the global LDP OAM configuration (at the **[edit protocols ldp oam]** or **[edit logical-systems logical-system-name protocols ldp oam]** hierarchy level) or in the configuration for the specified FEC (at the **[edit protocols ldp oam fec address]** or **[edit logical-systems logical-system-name protocols ldp oam fec address]** hierarchy level). Otherwise, the commit operation fails.

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

You can configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.

You can configure one of the following failure action options for the **failure-action** statement in the event of a BFD session failure on the LDP LSP:

- **remove-nexthop**—Removes the route corresponding to the next hop of the LSP's route at the ingress node when a BFD session failure event is detected.
- **remove-route**—Removes the route corresponding to the LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

To configure a failure action in the event of a BFD session failure on an LDP LSP, include either the **remove-nexthop** option or the **remove-route** option for the **failure-action** statement:

```
failure-action {  
    remove-nexthop;  
    remove-route;  
}
```

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

Configuring the Holddown Interval for the BFD Session

You can specify the duration the BFD session should be up before adding a route or next hop by configuring the **holddown-interval** statement at either the **[edit protocols ldp oam bfd-liveness-detection]** hierarchy level or at the **[edit protocols ldp oam fec address bfd-liveness-detection]** hierarchy level. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

```
holddown-interval seconds;
```

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

CHAPTER 8

Configuring Graceful Restart

- [LDP Graceful Restart on page 31](#)
- [Configuring LDP Graceful Restart on page 32](#)

LDP Graceful Restart

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

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

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

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

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

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

Distributing these other messages prevents the router from distributing incomplete information.

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

Configuring LDP Graceful Restart

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

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

To configure LDP graceful restart, see the following sections:

- [Enabling Graceful Restart on page 32](#)
- [Disabling LDP Graceful Restart or Helper Mode on page 32](#)
- [Configuring Reconnect Time on page 33](#)
- [Configuring Recovery Time and Maximum Recovery Time on page 34](#)

Enabling Graceful Restart

To enable LDP graceful restart, you also need to enable graceful restart on the router. To enable graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

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

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

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

Disabling LDP Graceful Restart or Helper Mode

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

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



```
}
}
```

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

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

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

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

The following LDP graceful restart configurations are possible:

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

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

Configuring Reconnect Time

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

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

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

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

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

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

Configuring Recovery Time and Maximum Recovery Time

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

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

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

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

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

CHAPTER 9

Configuring LDP Filtering and Apply Policies

- [Filtering Inbound LDP Label Bindings on page 35](#)
- [Filtering Outbound LDP Label Bindings on page 37](#)
- [Configuring OAM Ingress Policies for LDP on page 39](#)
- [Firewall Filter Match Conditions for MPLS Traffic on page 39](#)
- [Configuring the Prefixes Advertised into LDP from the Routing Table on page 41](#)

Filtering Inbound LDP Label Bindings

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

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

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

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. [Table 3 on page 35](#) lists the only **from** operators that apply to LDP received-label filtering.

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

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

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

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

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

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

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

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

Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

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

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

```

[edit]
protocols {
  ldp {
    import nosy-neighbor;
    ...
  }
}
policy-options {
  policy-statement nosy-neighbor {
    term first {
      from {
        neighbor 10.10.255.2;
        route-filter 131.108.0.0/16 orlonger accept;
        route-filter 0.0.0.0/0 orlonger reject;
      }
    }
    then accept;
  }
}

```

Filtering Outbound LDP Label Bindings

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

export [*policy-name*];

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

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in [Table 4 on page 37](#).

Table 4: to Operators for LDP Outbound-Label Filtering

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

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

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

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

Filtered labels are marked in the database:

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

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

Examples: Filtering Outbound LDP Label Bindings

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

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

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

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

```

    }
    then accept;
  }
  term block-the-rest {
    to {
      neighbor 10.10.255.2;
    }
    then reject;
  }
  then accept;
}
}

```

Configuring OAM Ingress Policies for LDP

Using the **ingress-policy** statement, you can configure an Operation, Administration, and Management (OAM) policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under **[edit protocols ldp oam bfd-liveness-detection]** are applied.

You configure the OAM ingress policy at the **[edit policy-options]** hierarchy level. To configure an OAM ingress policy, include the **ingress-policy** statement:

```
ingress-policy ingress-policy-name;
```

You can configure this statement at the following hierarchy levels:

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

Firewall Filter Match Conditions for MPLS Traffic

You can configure a firewall filter with match conditions for MPLS traffic (**family mpls**).



NOTE: The input-list *filter-names* and output-list *filter-names* statements for firewall filters for the mpls protocol family are supported on all interfaces with the exception of management interfaces and internal Ethernet interfaces (fxp or em0), loopback interfaces (lo0), and USB modem interfaces (umd).

Table 5 on page 39 describes the *match-conditions* you can configure at the **[edit firewall family mpls filter *filter-name* term *term-name* from]** hierarchy level.

Table 5: Firewall Filter Match Conditions for MPLS Traffic

Match Condition	Description
apply-groups	Specify which groups to inherit configuration data from. You can specify more than one group name. You must list them in order of inheritance priority. The configuration data in the first group takes priority over the data in subsequent groups.

Table 5: Firewall Filter Match Conditions for MPLS Traffic (*continued*)

Match Condition	Description
apply-groups-except	Specify which groups not to inherit configuration data from. You can specify more than one group name.
exp <i>number</i>	Experimental (EXP) bit number or range of bit numbers in the MPLS header. For <i>number</i> , you can specify one or more values from 0 through 7 in decimal, binary, or hexadecimal format. NOTE: This match condition is not supported on PTX series packet transport routers.
exp-except <i>number</i>	Do not match on the EXP bit number or range of bit numbers in the MPLS header. For <i>number</i> , you can specify one or more values from 0 through 7. NOTE: This match condition is not supported on PTX series packet transport routers.
forwarding-class <i>class</i>	Forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
forwarding-class-except <i>class</i>	Do not match on the forwarding class. Specify assured-forwarding , best-effort , expedited-forwarding , or network-control .
interface <i>interface-name</i>	Interface on which the packet was received. You can configure a match condition that matches packets based on the interface on which they were received. NOTE: If you configure this match condition with an interface that does not exist, the term does not match any packet.
interface-set <i>interface-set-name</i>	Match the interface on which the packet was received to the specified interface set. To define an interface set, include the interface-set statement at the [edit firewall] hierarchy level. NOTE: This match condition is not supported on PTX series packet transport routers. For more information, see <i>Filtering Packets Received on an Interface Set Overview</i> .
ip-version <i>number</i>	(Interfaces on Enhanced Scaling flexible PIC concentrators [FPCs] on supported T Series routers only) Inner IP version. To match MPLS-tagged IPv4 packets, match on the text synonym ipv4 . NOTE: This match condition is not supported on PTX series packet transport routers.
loss-priority <i>level</i>	Match the packet loss priority (PLP) level. Specify a single level or multiple levels: low , medium-low , medium-high , or high . Supported on M120 and M320 routers; M7i and M10i routers with the Enhanced CFEB (CFEB-E); and MX Series routers and EX Series switches. For IP traffic on M320, MX Series, and T Series routers with Enhanced II Flexible PIC Concentrators (FPCs), and EX Series switches, you must include the tri-color statement at the [edit class-of-service] hierarchy level to commit a PLP configuration with any of the four levels specified. If the tri-color statement is not enabled, you can only configure the high and low levels. This applies to all protocol families. For information about the tri-color statement, see <i>Configuring Tricolor Marking</i> . For information about using behavior aggregate (BA) classifiers to set the PLP level of incoming packets, see <i>Forwarding Classes Overview</i> .

Table 5: Firewall Filter Match Conditions for MPLS Traffic (*continued*)

Match Condition	Description
loss-priority-except level	Do not match the PLP level. For details, see the loss-priority match condition. NOTE: This match condition is not supported on PTX series packet transport routers.
Related Documentation	<ul style="list-style-type: none"> • <i>Guidelines for Configuring Firewall Filters</i> • <i>Firewall Filter Terminating Actions</i> • <i>Firewall Filter Nonterminating Actions</i>

Configuring the Prefixes Advertised into LDP from the Routing Table

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

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

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



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

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

Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

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

```
}  
}
```

Configuring LDP Traceroute and Collecting Statistics

- [Configuring LDP LSP Traceroute on page 43](#)
- [Collecting LDP Statistics on page 44](#)
- [Tracing LDP Protocol Traffic on page 46](#)

Configuring LDP LSP Traceroute

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

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

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

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

You can configure this statement at the following hierarchy levels:

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

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

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

Collecting LDP Statistics

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

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

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

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

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

This section includes the following topics:

- [LDP Statistics Output on page 44](#)
- [Disabling LDP Statistics on the Penultimate-Hop Router on page 45](#)
- [LDP Statistics Limitations on page 46](#)

LDP Statistics Output

The following sample output is from an LDP statistics file:

FEC	Type	Packets	Bytes	Shared
10.255.350.448/32	Transit	0	0	No
	Ingress	0	0	No

```

10.255.350.450/32  Transit      0      0      Yes
                   Ingress      0      0      No
10.255.350.451/32  Transit      0      0      No
                   Ingress      0      0      No
220.220.220.1/32   Transit      0      0      Yes
                   Ingress      0      0      No
220.220.220.2/32   Transit      0      0      Yes
                   Ingress      0      0      No
220.220.220.3/32   Transit      0      0      Yes
                   Ingress      0      0      No
May 28 15:02:05, read 12 statistics in 00:00:00 seconds

```

The LDP statistics file includes the following columns of data:

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

Disabling LDP Statistics on the Penultimate-Hop Router

Gathering LDP traffic statistics at the penultimate-hop router can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement. For routers reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```

traffic-statistics {
  no-penultimate-hop;
}

```

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



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

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

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

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

LDP Statistics Limitations

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

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

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

Tracing LDP Protocol Traffic

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

- [Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 46](#)
- [Tracing LDP Protocol Traffic Within FECs on page 47](#)
- [Examples: Tracing LDP Protocol Traffic on page 48](#)

Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels

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

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

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

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

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

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

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

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

Tracing LDP Protocol Traffic Within FECs

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

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

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

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

This feature has the following limitations:

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

Examples: Tracing LDP Protocol Traffic

Trace LDP path messages in detail:

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

Trace all LDP outgoing messages:

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

Trace all LDP error conditions:

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

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

```
[edit]
protocols {
  ldp {
```



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

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

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


PART 3

Configuration Statements and Operational Commands

- Configuration Statements on page 53
- Operational Commands on page 113

CHAPTER 11

Configuration Statements

- [\[edit protocols bgp\] Hierarchy Level on page 54](#)
- [\[edit protocols ldp\] Hierarchy Level on page 61](#)
- [\[edit protocols mpls\] Hierarchy Level on page 63](#)
- [allow-subnet-mismatch on page 64](#)
- [authentication-algorithm on page 65](#)
- [authentication-key \(Protocols LDP\) on page 67](#)
- [bfd-liveness-detection \(Protocols LDP\) on page 68](#)
- [deaggregate on page 69](#)
- [disable \(Protocols LDP\) on page 70](#)
- [dod-request-policy on page 71](#)
- [downstream-on-demand on page 71](#)
- [ecmp on page 72](#)
- [egress-policy on page 72](#)
- [explicit-null \(Protocols LDP\) on page 73](#)
- [export \(Protocols LDP\) on page 73](#)
- [failure-action \(Protocols LDP\) on page 74](#)
- [fec on page 75](#)
- [graceful-restart \(Protocols LDP\) on page 76](#)
- [hello-interval \(Protocols LDP\) on page 77](#)
- [helper-disable \(LDP\) on page 78](#)
- [holddown-interval on page 79](#)
- [hold-time \(Protocols LDP\) on page 80](#)
- [igp-synchronization on page 81](#)
- [import \(Protocols LDP\) on page 82](#)
- [ingress-policy on page 83](#)
- [interface \(Protocols LDP\) on page 84](#)
- [keepalive-interval on page 85](#)
- [keepalive-timeout on page 86](#)

- [l2-smart-policy](#) on page 86
- [label-withdrawal-delay](#) on page 87
- [ldp](#) on page 88
- [ldp-p2mp](#) on page 91
- [log-updown \(Protocols LDP\)](#) on page 92
- [make-before-break \(LDP\)](#) on page 93
- [maximum-neighbor-recovery-time](#) on page 94
- [no-forwarding](#) on page 95
- [oam \(Protocols LDP\)](#) on page 96
- [p2mp \(Protocols LDP\)](#) on page 97
- [periodic-traceroute](#) on page 98
- [policing \(Protocols LDP\)](#) on page 100
- [preference \(Protocols LDP\)](#) on page 101
- [reconnect-time](#) on page 102
- [recovery-time](#) on page 103
- [session \(ldp\)](#) on page 104
- [session-protection](#) on page 105
- [strict-targeted-hellos](#) on page 105
- [targeted-hello](#) on page 106
- [traceoptions \(Protocols LDP\)](#) on page 107
- [track-igp-metric](#) on page 109
- [traffic-statistics \(Protocols LDP\)](#) on page 110
- [transport-address](#) on page 112

[edit protocols bgp] Hierarchy Level

Several statements in the **[edit protocols mpls]** hierarchy are valid at numerous locations within it. To make the complete hierarchy easier to read, the repeated statements are listed in “[Common BGP Family Options](#)” on page 54 and that section is referenced at the appropriate locations in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 55.

- [Common BGP Family Options](#) on page 54
- [Complete \[edit protocols bgp\] Hierarchy](#) on page 55

Common BGP Family Options

This section lists statements that are valid at the following hierarchy levels, and is referenced at those levels in “[Complete \[edit protocols bgp\] Hierarchy](#)” on page 55 instead of the statements being repeated.

- **[edit protocols bgp family inet (any | flow | labeled-unicast | multicast | unicast)]**
- **[edit protocols bgp family inet6 (any | labeled-unicast | multicast | unicast)]**

- [edit protocols bgp family (evpn | inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) signaling]
- [edit protocols bgp family inet-vpn (any | flow | multicast | unicast)]
- [edit protocols bgp family inet6-vpn (any | multicast | unicast)]
- [edit protocols bgp family iso-vpn unicast]

The common BGP family options are as follows:

```
accepted-prefix-limit {
  maximum number;
  teardown <percentage> <idle-timeout (forever | minutes)>;
}
damping;
loops number;
prefix-limit {
  maximum number;
  teardown <percentage> <idle-timeout (forever | minutes)>;
}
rib-group group-name;
topology name {
  community {
    target identifier;
  }
}
```

Complete [edit protocols bgp] Hierarchy

The statement hierarchy listed in this section can also be included at the [edit logical-systems *logical-system-name*] hierarchy level.

```
protocols {
  bgp {
    disable;
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-from-main-vpn-tables;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain key-chain;
    bfd-liveness-detection {
      authentication {
        algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 |
          meticulous-keyed-sha-1 | simple-password);
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      holddown-interval milliseconds;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
    }
  }
}
```

```

no-adaptation;
session-mode (automatic | multihop | single-hop);
transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
}
version (1 | automatic);
}
bmp {
    monitor (disable | enable);
    route-monitoring {
        none;
        post-policy {
            exclude-non-eligible;
        }
        pre-policy {
            exclude-non-feasible;
        }
    }
}
cluster cluster-identifier;
damping;
description text-description;
export [ policy-names ];
family family-name {
    ... the family subhierarchies appear after the main [edit protocols bgp] hierarchy ...
}
unconfigured-peer-graceful-restart;
graceful-restart {
    disable;
    restart-time seconds;
    stale-routes-time seconds;
}
graceful-restart {
    long-lived {
        receiver {
            enable;
            disable;
        }
        advertise-to-non-llgr-neighbor {
            omit-no-export;
        }
    }
}
graceful-restart {
    disable-notification-flag;
    disable-notification-extensions {
        omit-no-export;
    }
}
forwarding-state-bit (from-fib | set); /* Configurable to be common for all address
    families */
forwarding-state-bit (as-rr-client | from-fib); /* Configurable for each address family
    */
long-lived {
    restarter {
        disable;
    }
}

```



```

        stale-time interval;
    }
}
group group-name {
    ... the group subhierarchy appears after the main [edit protocols bgp] hierarchy ...
}
hold-time seconds;
idle-after-switch-over (seconds | forever);
import [ policy-names ];
include-mp-next-hop;
ipsec-sa ipsec-sa;
keep (all | none);
local-address address;
local-as autonomous-system <loops number> <alias> <private>;
local-interface interface-name;
local-preference local-preference;
log-updown;
metric-out (metric | igp (delay-med-update | offset) | minimum-igp offset);
mtu-discovery;
multihop {
    no-nexthop-change;
    ttl tvl-value;
}
no-aggregator-id;
no-client-reflect;
out-delay seconds;
outbound-route-filter {
    bgp-orf-cisco-mode;
    prefix-based {
        accept {
            inet;
            inet6;
        }
    }
}
passive;
path-selection {
    always-compare-med;
    as-path-ignore;
    cisco-non-deterministic;
    external-router-id;
    med-plus-igp {
        igp-multiplier number;
        med-multiplier number;
    }
}
peer-as autonomous-system;
preference preference;
remove-private;
tcp-mss segment-size;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
        no-world-readable>;
    flag flag <flag-modifier> <disable>;
}

```

```

    tcp-aggressive-transmission;
    vpn-apply-export;
}

bgp {
  family inet {
    (any | multicast) {
      ... statements in Common BGP Family Options on page 54 ...
    }
    flow {
      ... statements in Common BGP Family Options on page 54 PLUS ...
      no-validate [ validation-procedure-names ];
    }
    labeled-unicast {
      ... statements in Common BGP Family Options on page 54 PLUS ...
      add-path {
        receive;
        send {
          path-count number;
          prefix-policy [ policy-names ];
        }
      }
      aggregate-label {
        community community-name;
      }
      aigp [disable];
      explicit-null connected-only;
      per-group-label;
      per-prefix-label;
      protection;
      resolve-vpn;
      rib (inet.3 | inet6.3);
      traffic-statistics {
        file filename <files number> <size maximum-file-size> <world-readable |
          no-world-readable>;
        interval seconds;
      }
    }
  }
  unicast {
    ... statements in Common BGP Family Options on page 54 PLUS ...
    add-path {
      receive;
      send {
        path-count number;
        prefix-policy [ policy-names ];
      }
    }
    topology name {
      community target identifier;
    }
  }
}

bgp {
  family inet6 {

```

```

    (any | multicast) {
        ... statements in Common BGP Family Options on page 54 ...
    }
    labeled-unicast {
        ... statements in Common BGP Family Options on page 54 PLUS ...
        add-path {
            receive;
            send {
                path-count number;
                prefix-policy [ policy-names ];
            }
        }
        aggregate-label {
            community community-name;
        }
        aigp [disable];
        explicit-null;
        per-group-label;
        protection;
        traffic-statistics {
            file filename <files number> <size maximum-file-size> <world-readable |
                no-world-readable>;
            interval seconds;
        }
    }
    unicast {
        ... statements in Common BGP Family Options on page 54 PLUS ...
        topology name {
            community target identifier;
        }
    }
}

bgp {
    family (evpn | inet-mdt | inet-mvpn | inet6-mvpn | l2vpn) {
        auto-discovery-only; # for l2vpn
        signaling {
            ... statements in Common BGP Family Options on page 54 ...
        }
    }
}

bgp {
    family inet-vpn {
        (any | multicast | unicast) {
            ... statements in Common BGP Family Options on page 54 PLUS ...
            aggregate-label <community community-name>;
        }
        flow {
            ... statements in Common BGP Family Options on page 54 ...
        }
    }
}

bgp {

```

```

family inet6-vpn {
  (any | multicast | unicast) {
    ... statements in Common BGP Family Options on page 54 PLUS ...
    aggregate-label <community community-name>;
  }
}

bgp {
  family iso-vpn {
    unicast {
      ... statements in Common BGP Family Options on page 54 PLUS ...
      aggregate-label <community community-name>;
    }
  }
}

bgp {
  family route-target {
    accepted-prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    proxy-generate <route-target-policy route-target-policy-name>;
  }
}

bgp {
  group group-name {
    ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
    allow [ all ip-prefix</prefix-length> ];
    as-override;
    multipath <multiple-as>;
    neighbor address {
      ... the neighbor subhierarchy appears after the main [edit protocols bgp group
        group-name] hierarchy ...
    }
    type (external | internal);
    ... BUT NOT ...
    disable; # NOT valid at this level
    group group-name { ... } # NOT valid at this level
    path-selection { ... } # NOT valid at this level
  }

  group group-name {
    neighbor address {
      ... same statements as at the [edit protocols bgp] hierarchy level PLUS ...
      as-override;
      multipath <multiple-as>;
      ... BUT NOT ...

```

```

        disable; # NOT valid at this level
        group group-name { ... } # NOT valid at this level
        neighbor address { ... } # NOT valid at this level
        path-selection { ... } # NOT valid at this level
    }
}
}
}

```

**Related
Documentation**

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols ldp] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
  ldp {
    (deaggregate | no-deaggregate);
    dod-request-policy [ policy-names ];
    egress-policy [ policy-names ];
    explicit-null;
    export [ policy-names ];
    graceful-restart {
      disable;
      helper-disable;
      maximum-neighbor-reconnect-time seconds;
      maximum-neighbor-recovery-time seconds;
      reconnect-time seconds;
      recovery-time seconds;
    }
    igp-synchronization holddown-interval seconds;
    import [ policy-names ];
    interface interface-name {
      (allow-subnet-mismatch | no-allow-subnet-mismatch);
      disable;
      hello-interval seconds;
      hold-time seconds;
      transport-address (interface | router-id);
    }
    keepalive-interval seconds;
    keepalive-timeout seconds;
    l2-smart-policy;
    log-updown {
      trap disable;
    }
    next-hop {
      merged {
        policy [ policy-names ];
      }
    }
    no-forwarding;
    oam {

```

```

... the oam subhierarchy appears after the main [edit protocols ldp] hierarchy ...
}
p2mp {
  root-address root-address {
    lsp-id id;
  }
}
policing {
  fec class-address {
    ingress-traffic filter-name;
    transit-traffic filter-name;
  }
}
preference preference;
session destination-address {
  authentication-algorithm algorithm;
  authentication-key key;
  authentication-key-chain key-chain;
  downstream-on-demand;
}
session-protection <timeout seconds>;
strict-targeted-hellos;
targeted-hello {
  hello-interval seconds;
  hold-time seconds;
}
traceoptions {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
track-igp-metric;
traffic-statistics {
  file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
  interval seconds;
  no-penultimate-hop;
}
transport-address (address | interface | router-id);
}

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;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
}

```

```

    }
    version (1 | automatic);
  }
  fec class-address {
    bfd-liveness-detection {
      ... same statements as at the [edit protocols ldp oam bfd-liveness-detection]
        hierarchy level ...
    }
    no-bfd-liveness-detection;
    periodic-traceroute {
      ... same statements as at the [edit protocols ldp oam periodic-traceroute]
        hierarchy level PLUS ...
    }
    disable;
  }
}
ingress-policy [ policy-names ];
periodic-traceroute {
  exp cos-value;
  fanout next-hops;
  frequency minutes;
  paths number;
  retries number;
  source address;
  ttl number;
  wait seconds;
}
}
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

[edit protocols mpls] Hierarchy Level

- Complete [edit protocols mpls] Hierarchy on page 63

Complete [edit protocols mpls] Hierarchy

The statement hierarchy listed in this section can also be included at the **[edit logical-systems logical-system-name]** hierarchy level.

```

protocols {
  mpls {
    disable;
    interface (interface-name | all) {
      always-mark-connection-protection-tlv;
      disable;
      admin-group [ group-names ];
      srlg srlg-name;
      static {
        protection-revert-time seconds;
      }
    }
    switch-away-lsps;
  }
}

```

```

    }
    egress-protection {
        context-identifier context-id {
            primary | protector;
            metric igp-metric-value;
            advertise-mode (stub-alias | stub-proxy);
        }
    }
    ipv6-tunneling;
    priority setup-priority hold-priority;
    traceoptions {
        file filename <files number> <size maximum-file-size> <world-readable |
            no-world-readable>;
        flag flag;
    }
}

```

- Related Documentation**
- *Notational Conventions Used in Junos OS Configuration Hierarchies*
 - *[edit protocols] Hierarchy Level*

allow-subnet-mismatch

Syntax	allow-subnet-mismatch;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit protocols ldp interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.
Default	The source address in the LDP link hello packet is matched against the interface address.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Miscellaneous LDP Properties</i>

authentication-algorithm

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



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

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

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

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

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

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



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

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

Related Documentation

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

authentication-key (Protocols LDP)

Syntax `authentication-key md5-authentication-key;`

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

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

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

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

Related Documentation

- *Configuring Miscellaneous LDP Properties*

bfd-liveness-detection (Protocols LDP)

Syntax	<pre>bfd-liveness-detection { detection-time threshold <i>milliseconds</i>; ecmp; failure-action { remove-nexthop; remove-route; } holddown-interval <i>seconds</i>; minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; minimum-transmit-interval <i>milliseconds</i>; multiplier <i>detection-time-multiplier</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (0 1 automatic); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec address], [edit protocols ldp oam], [edit protocols ldp oam fec address]
Release Information	Statement introduced in Junos OS Release 7.6. Support for the bfd-liveness-detection statement at the [edit protocols ldp oam fec address] hierarchy level and the ecmp option added in Junos OS Release 9.0. Support for the failure-action statement with the remove-nexthop and remove-route options and the holddown-interval statement added in Junos OS Release 9.4.
Description	Enable Bidirectional Forwarding Detection (BFD) for all MPLS LSPs or for just a specific LSP.
Options	<p>minimum-interval—Minimum transmit and receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-receive-interval—Minimum receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-transmit-interval—Minimum transmit interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>multiplier—Detection time multiplier. Range: 50 through 255</p>

Default: 3

The other options are explained separately.

Required Privilege Level	routing—To view this statement in the configuration.
	routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring BFD for LDP LSPs on page 25

deaggregate

Syntax	deaggregate no-deaggregate;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Control forwarding equivalence class (FEC) deaggregation on the router. The use of the deaggregate statement in LDP is a standard practice that we recommend for LDP deployments.
Default	Deaggregation is disabled on the router.
Options	deaggregate —Deaggregate FECs. no-deaggregate —Aggregate FECs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring FEC Deaggregation on page 21

disable (Protocols LDP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit protocols ldp graceful-restart], [edit protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.
Default	LDP is enabled on interfaces configured with the LDP interface statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the [edit routing-options] hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling and Disabling LDP on page 16• Configuring LDP Graceful Restart on page 32

dod-request-policy

Syntax	<code>dod-request-policy <i>dod-request-policy-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the name of the LDP downstream on demand request policy. LDP sends label request messages only for those FECs matching in the downstream on demand request policy.
Options	<i>dod-request-policy-name</i> —Specify the name of the downstream on demand request policy.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring LDP Downstream on Demand</i>

downstream-on-demand

Syntax	<code>downstream-on-demand;</code>
Hierarchy Level	[edit logical systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>], [edit protocols ldp session <i>session-address</i>]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring LDP Downstream on Demand</i>

ecmp

Syntax	<code>ecmp;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the ecmp statement, you must also configure the periodic-traceroute statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the periodic-traceroute statement at the global hierarchy level ([edit protocols ldp oam]) while only configuring the ecmp statement for a specific FEC ([edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring ECMP-Aware BFD for LDP LSPs on page 28

egress-policy

Syntax	<code>egress-policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Control the prefixes advertised into LDP.
Default	Only the loopback address is advertised.
Options	<i>policy-names</i> —Name of one or more routing policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Prefixes Advertised into LDP from the Routing Table on page 41

explicit-null (Protocols LDP)

Syntax	<code>explicit-null;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Advertise label 0 to the egress router of a label-switched path (LSP).
Default	If you do not include the explicit-null statement in the MPLS configuration, label 3 (implicit null) is advertised.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Miscellaneous LDP Properties

export (Protocols LDP)

Syntax	<code>export [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.
Options	<i>policy-names</i> —Name of one or more routing policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Filtering Outbound LDP Label Bindings on page 37

failure-action (Protocols LDP)

Syntax	<pre>failure-action { remove-nexthop; remove-route; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-livenessss-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-livenessss-detection], [edit protocols ldp oam bfd-livenessss-detection], [edit protocols ldp oam fec <i>address</i> bfd-livenessss-detection]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.
Options	<p>remove-nexthop—Remove a route corresponding to a next hop of the LSP's route at the ingress node when a BFD session failure event is detected.</p> <p>remove-route—Remove the route corresponding to an LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Failure Action for the BFD Session on an LDP LSP on page 28

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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).
Options	<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 on page 25](#)

graceful-restart (Protocols LDP)

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



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

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

Related Documentation • [Configuring LDP Graceful Restart on page 32](#)

hello-interval (Protocols LDP)

Syntax	<code>hello-interval <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello],</p> <p>[edit protocols ldp interface <i>interface-name</i>],</p> <p>[edit protocols ldp targeted-hello],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for LDP targeted hellos added in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the hello-interval statement.
Options	<p><i>seconds</i>—Length of time between transmission of hello packets.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 5 seconds for link hello messages, 15 seconds for targeted hello messages</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the LDP Timer for Hello Messages on page 17

helper-disable (LDP)

Syntax	helper-disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.
Default	Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LDP Graceful Restart on page 32

holddown-interval

Syntax	<code>holddown-interval <i>holddown-interval</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-livenesss-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-livenesss-detection], [edit protocols ldp oam bfd-livenesss-detection], [edit protocols ldp oam fec <i>address</i> bfd-livenesss-detection]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Specify how long the BFD session should be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.
Options	<i>holddown-interval</i> —Number of seconds the BFD session should remain up before adding the route or next hop. Default: 0 seconds Range: 0 through 65,535 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Holddown Interval for the BFD Session on page 29

hold-time (Protocols LDP)

Syntax	<code>hold-time seconds;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ldp targeted-hello],</code> <code>[edit protocols ldp interface <i>interface-name</i>],</code> <code>[edit protocols ldp targeted-hello],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Support for LDP targeted hellos added in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the hold-time statement.
Options	seconds —Hold-time value. Range: 1 through 65,535 seconds Default: 15 seconds for link hello messages, 45 seconds for targeted hello messages
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Delay Before LDP Neighbors Are Considered Down on page 18

igp-synchronization

Syntax	<code>igp-synchronization holddown-interval <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.
Options	holddown-interval <i>seconds</i> —Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. Default: 10 seconds Range: 10 through 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Miscellaneous LDP Properties</i>

import (Protocols LDP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.
Options	<i>policy-names</i> —Name of one or more routing policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Filtering Inbound LDP Label Bindings on page 35

ingress-policy

Syntax	<code>ingress-policy [<i>ingress-policy-names</i>];</code>
Hierarchy Level	[edit logical-system <i>logical-system-name</i> protocols ldp entropy-label], [edit logical-system <i>logical-system-name</i> protocols ldp oam], [edit protocols ldp entropy-label], [edit protocols ldp oam]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced at the [edit protocols ldp entropy-label] hierarchy level in Junos OS Release 14.1.
Description	<p>Configure an LDP ingress policy for either the entropy label or Operation, Administration, and Management (OAM).</p> <p>For OAM, configure the ingress policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under [edit protocols ldp oam bfd-liveness-detection] are applied.</p>
Options	<i>ingress-policy-names</i> —Specify the names of the ingress policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring OAM Ingress Policies for LDP on page 39 • Configuring the Entropy Label for LSPs

interface (Protocols LDP)

Syntax	<pre>interface <i>interface-name</i> { disable; hello-interval <i>seconds</i>; hold-time <i>seconds</i>; transport-address (interface loopback); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable LDP on one or more router interfaces.
Default	LDP is disabled on all interfaces.
Options	<i>interface-name</i> —Name of an interface. To configure all interfaces, specify all . The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling and Disabling LDP on page 16

keepalive-interval

Syntax	<code>keepalive-interval <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the keepalive interval value.
Options	<i>seconds</i> —Keepalive value. Range: 1 through 65,535 Default: 10 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Interval for LDP Keepalive Messages on page 19

keepalive-timeout

Syntax	keepalive-timeout <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.
Options	seconds —Keepalive timeout value. Range: 1 through 65,535 Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the LDP Keepalive Timeout on page 19

l2-smart-policy

Syntax	l2-smart-policy;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LDP IPv4 FEC Filtering on page 23

label-withdrawal-delay

Syntax	label-withdrawal-delay <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Delay the withdrawal of labels to reduce router workload during IGP convergence.
Options	<i>seconds</i> —Configure the number of seconds to wait before withdrawing labels for the LDP LSPs. Default: 60 seconds Range: 0 through 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Miscellaneous LDP Properties</i>

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Enable LDP routing on the router or switch.

You must include the **ldp** statement in the configuration to enable LDP on the router or switch.

Default LDP is disabled on the router.

Options The other statements are explained separately.

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

Related Documentation

- [Minimum LDP Configuration on page 15](#)
- [Enabling and Disabling LDP on page 16](#)

ldp-p2mp

Syntax	ldp-p2mp;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> wildcard-source],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> source <i>source-prefix</i>],</p> <p>[edit protocols mvpn inter-region-template template <i>template-name</i> all-regions],</p> <p>[edit protocols mvpn inter-region-template template <i>template-name</i> region <i>region-name</i>],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet wildcard-source],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective wildcard-group-inet6 wildcard-source],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> wildcard-source],</p> <p>[edit routing-instances <i>instance-name</i> provider-tunnel selective group <i>group-prefix</i> source <i>source-prefix</i>]</p>
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify a point-to-multipoint provider tunnel with LDP signalling for an MBGP MVPN.
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 Point-to-Multipoint LDPLSPs as the Data Plane for Intra-AS MBGP MVPNs

log-updown (Protocols LDP)

Syntax	log-updown { trap disable; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Disable LDP traps on the router, logical system, or routing instance.
Options	trap disable —Disable LDP traps. Default: LDP traps are enabled on the router.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Miscellaneous LDP Properties</i>

make-before-break (LDP)

Syntax	<pre>make-before-break { timeout <i>seconds</i>; switchover-delay <i>seconds</i>; }</pre>
Hierarchy Level	[edit protocols ldp]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Configures make before break (MBB) for multicast LDP (MLDP) link protection to ensure minimum packet loss when attempting to signal a new label-switched path (LSP) before tearing down the old LSP path.
Options	<p>timeout <i>seconds</i>—Specify a value to change a make -before-break timeout for point-to-multipoint LSPs. Even if an MBB acknowledgment is not received for a point-to-multipoint LSP before the specified timeout period expires, the label-switching router (LSR) performs an MBB switchover from the old LSR to the new upstream LSR.</p> <p>Range: 1 through 300 seconds</p> <p>Default: 30 seconds</p> <p>switchover-delay <i>seconds</i>—Specify a value to change switchover delay for a point-to-multipoint LSP from the old LSR to the new upstream LSR. If an MBB acknowledgment is received on a point of local repair (PLR) router, the PLR waits for the specified seconds to switch its upstream LSR from the old LSR to the new LSR.</p> <p>Range: 1 through 300 seconds</p> <p>Default: 30 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring LDP Link Protection</i>

maximum-neighbor-recovery-time

Syntax	<code>maximum-neighbor-recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement changed from maximum-recovery-time to maximum-neighbor-recovery-time in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
Options	seconds —Configure the maximum recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Recovery Time and Maximum Recovery Time on page 34• <i>Configuring Graceful Restart Options for LDP</i>• <i>no-strict-lsa-checking</i>• <i>recovery-time</i>

no-forwarding

Syntax	no-forwarding;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Do not add ingress routes to the inet.0 routing table even if traffic-engineering bgp-igp (configured at the [edit protocols mpls] hierarchy level) is enabled.
Default	The no-forwarding statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when traffic-engineering bgp-igp is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Miscellaneous LDP Properties</i> • <i>Configuring Virtual-Router Routing Instances in VPNs</i>

oam (Protocols LDP)

```
Syntax  oam {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
            holddown-interval milliseconds;
            ingress-policy ingress-policy-name;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        fec fec-address;
        ingress-policy ingress-policy-name;
        lsp-ping-interval seconds;
        periodic-traceroute {
            disable;
            exp exp-value;
            fanout fanout-value;
            frequency minutes;
            paths number-of-paths;
            retries retry-attempts;
            source address;
            ttl ttl-value;
            wait seconds;
        }
    }
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols *ldp*]
[edit protocols *ldp*]

Release Information Statement introduced in Junos OS Release 7.6.
lsp-ping-interval option introduced in Junos OS Release 9.4.

Description Configure Operation, Administration, and Maintenance (OAM) and Bidirectional Forwarding Detection (BFD) protocol for LDP.

Options ***fec fec-address***—Specify the forwarding equivalence class (FEC) address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.

lsp-ping-interval *seconds*—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command.

Default: 60 seconds

Range: 30 through 3,600 seconds

The remaining statements are explained separately.

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

Related Documentation

- [Configuring BFD for LDP LSPs on page 25](#)

p2mp (Protocols LDP)

Syntax p2mp{
 root-address *root-address*{
 lsp-id *id*;
 }

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

Release Information Statement introduced in Junos OS Release 11.2.

Description Enable point-to-multipoint MPLS LSPs in an LDP-signaled LSP.

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

Related Documentation

- *Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs*
- *Point-to-Multipoint LSPs Overview*

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">• Configuring LDP LSP Traceroute on page 43
------------------------------	---

policing (Protocols LDP)

Syntax	<pre>policing { fec <i>fec-address</i> { ingress-traffic <i>filter-name</i>; transit-traffic <i>filter-name</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable policing of forwarding equivalence classes (FECs) for LDP.
Options	<p>fec <i>fec-address</i>—Specify the address for the FEC.</p> <p>ingress-traffic <i>filter-name</i>—Specify the name of the filter for policing ingress FEC traffic.</p> <p>transit-traffic <i>filter-name</i>—Specify the name of the filter for policing transit FEC traffic.</p>
Required Privilege Level	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 Policers for LDP FECs on page 22

preference (Protocols LDP)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the route preference level for LDP routes.
Options	<i>preference</i> —Preferred value. Range: 0 through 255 Default: 9
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LDP Route Preferences on page 20

reconnect-time

Syntax	<code>reconnect-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
Options	seconds —Time required for reconnection. Range: 30 through 300 Default: 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LDP Graceful Restart on page 32 on <i>MPLS Applications Feature Guide for Routing Devices</i>• <i>Configuring Graceful Restart Options for LDP</i>

recovery-time

Syntax	<code>recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the amount of time a router waits for LDP to restart gracefully.
Options	seconds —Configure the recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Recovery Time and Maximum Recovery Time on page 34

session (ldp)

Syntax	<pre>session address { authentication-algorithm <i>algorithm</i>; authentication-key <i>authentication-key</i>; authentication-key-chain <i>key-chain-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. authentication-algorithm statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the address for the remote end of the LDP session. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Miscellaneous LDP Properties</i>

session-protection

Syntax	session-protection { timeout <i>seconds</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Description	Configure when an LDP session is torn down and resigaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.
Options	timeout <i>seconds</i> —Time in seconds before the LDP session is torn down and resigaled. Range: 1 through 65,535 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Miscellaneous LDP Properties

strict-targeted-hellos

Syntax	strict-targeted-hellos;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Strict Targeted Hello Messages for LDP on page 20

targeted-hello

Syntax	targeted-hello { hello-interval <i>seconds</i> ; hold-time <i>seconds</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the LDP timer and LDP hold time for targeted hellos.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the LDP Timer for Hello Messages on page 17• Configuring the Delay Before LDP Neighbors Are Considered Down on page 18

traceoptions (Protocols LDP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols <i>ldp</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>], [edit protocols <i>ldp</i>], [edit routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>match-on address option for the filter flag modifier added in Junos OS Release 10.4.</p> <p>nsr-synchronization and p2mp-nsr-synchronization operations for flag statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Specify LDP protocol-level trace options.
Default	The default LDP protocol-level trace options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory ldp-log. We recommend that you place LDP tracing output in the file ldp-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none"> • address—Operation of address and address withdrawal messages • binding—Label-binding operations • error—Error conditions • event—Protocol events

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

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

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

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

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

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

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

Default: 1 MB

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

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

Required Privilege Level	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing LDP Protocol Traffic on page 46 • <i>Network Management Administration Guide for Routing Devices</i>

track-igp-metric

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

traffic-statistics (Protocols LDP)

Syntax	<pre>traffic-statistics { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; interval <i>seconds</i>; no-penultimate-hop; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of LDP statistics files. When a statistics file named <i>ldp-stat</i> reaches its maximum size, it is renamed <i>ldp-stat.0</i>, then <i>ldp-stat.1</i>, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must include the size statement to specify the maximum file size.</p> <p>interval <i>seconds</i>—(Optional) Specify the interval at which the statistics are polled and written to the file.</p> <p>Default: 300 seconds (5 minutes)</p> <p>no-penultimate-hop—(Optional) Do not collect traffic statistics on the penultimate hop router.</p> <p>no-world-readable—(Optional) Prevent all users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named <i>ldp-stat</i> reaches this size, it is renamed <i>ldp-stat.0</i>. When <i>ldp-stat</i> again reaches this size, <i>ldp-stat.0</i> is renamed <i>ldp-stat.1</i> and <i>ldp-stat</i> is renamed <i>ldp-stat.0</i>. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten.</p> <p>Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB</p> <p>Range: 10 KB through the maximum file size supported on your system</p>

Default: 1 MB

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

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

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

Related Documentation	<ul style="list-style-type: none">• Collecting LDP Statistics on page 44
------------------------------	--

transport-address

Syntax	<code>transport-address (interface router-id);</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ldp],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>ldp],</code> <code>[edit protocols ldp],</code> <code>[edit protocols ldp interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enables you to configure the IP address used to specify the TCP session for the LDP session. Routers must first establish a TCP session between one another before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.
Default	router-id
Options	interface —The first IP address on the interface is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the interface option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the router-id option. router-id —The router identifier is used as the transport address. Unless otherwise configured, the router identifier is the loopback address.
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Specifying the Transport Address Used by LDP on page 16

CHAPTER 12

Operational Commands

- ping mpls ldp
- show ldp database
- show ldp session
- show ldp traffic-statistics
- show ldp session

ping mpls ldp

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

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

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

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

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

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

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

Required Privilege Level network

List of Sample Output [ping mpls ldp fec count on page 115](#)
[ping mpls ldp p2mp root-addr lsp-id on page 115](#)

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

Sample Output

ping mpls ldp fec count

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

ping mpls ldp p2mp root-addr lsp-id

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

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

show ldp database

Syntax	<pre>show ldp database <brief detail extensive> <inet l2circuit> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <p2mp> <session <i>session</i>> <summary></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>summary option introduced in Junos OS Release 14.2.</p>
Description	Display entries in the LDP database.
Options	<p>none—Display standard information about all entries in the LDP database for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>inet l2circuit—(Optional) Display only IPv4 or Layer 2 circuit bindings.</p> <p>instance <i>instance-name</i>—(Optional) Display routing instance information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>p2mp—(Optional) Display point-to-multipoint binding information.</p> <p>session <i>session</i>—(Optional) Display database for the specified session only. <i>session</i> is the destination address of the LDP session.</p> <p>summary—(Optional)—Display summary output. This option displays the number of labels received and advertised for each LDP session.</p>
Required Privilege Level	view
List of Sample Output	<p>show ldp database (master) on page 120</p> <p>show ldp database (standby) on page 121</p> <p>show ldp database l2circuit detail on page 121</p> <p>show ldp database l2circuit extensive on page 122</p> <p>show ldp database p2mp (master) on page 122</p> <p>show ldp database p2mp (standby) on page 122</p> <p>show ldp database p2mp (master) on page 123</p> <p>show ldp database p2mp (standby) on page 123</p> <p>show ldp database session on page 123</p> <p>show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 124</p>

[show ldp database \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 124](#)
[show ldp database summary on page 125](#)

Output Fields [Table 6 on page 118](#) describes the output fields for the **show ldp database** command. Output fields are listed in the approximate order in which they appear.

Table 6: show ldp database Output Fields

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

Table 6: show ldp database Output Fields (*continued*)

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

Table 6: show ldp database Output Fields (*continued*)

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

Sample Output

show ldp database (master)

```

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

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

```



```

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

```

show ldp database (standby)

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

```
Input label database, 10.255.107.236:0--10.255.107.234:0
```

```

Label Prefix
299808 10.255.107.230/32
State: Active
Age: 1d 2:46:36
Standby binding state:
Map messages: 1
Release messages: 0

```

```

Label Prefix
301136 10.255.107.232/32
State: Active
Age: 1d 2:46:36
Standby binding state:
Map messages: 1
Release messages: 0

```

```

Label Prefix
3      10.255.107.234/32
State: Active
Age: 1d 2:46:36
Standby binding state:
Map messages: 1
Release messages: 0

```

```

Label Prefix
302480 10.255.107.236/32
State: Active
Age: 1d 2:46:36
Standby binding state:
Map messages: 1
Release messages: 0

```

```
Output label database, 10.255.107.236:0--10.255.107.234:0
```

```

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

```

show ldp database l2circuit detail

```
user@host> show ldp database l2circuit detail
```

```

Input label database, 10.255.245.44:0--10.255.245.45:0
Label Prefix

```

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

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

show ldp database l2circuit extensive

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

show ldp database p2mp (master)

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

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

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

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

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

show ldp database p2mp (standby)

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

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

```

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

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

```

show ldp database p2mp (master)

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

```

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

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

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

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

```

show ldp database p2mp (standby)

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

```

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

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

```

show ldp database session

```

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

```

```

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

```

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

```

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

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

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

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

```

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

```

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

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

```

```

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

```

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

```

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

```

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

```

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

```

show ldp database summary

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

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

show ldp session

Syntax	<pre>show ldp session <brief detail extensive> <auto-targeted> <destination> <instance instance-name> <logical-system (all logical-system-name)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>auto-targeted option added in Junos OS Release 14.2.</p>
Description	Display information about Label Distribution Protocol (LDP) sessions.
Options	<p>none—Display standard information about all LDP sessions for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>auto-targeted—(Optional) Display information about LDP sessions that are automatically targeted using loopback addresses.</p> <p>destination—(Optional) Restrict LDP session display to the specified address.</p> <p>instance instance-name—(Optional) Display routing instance information for the specified instance. If instance-name is omitted, information is displayed for the master instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ldp session
List of Sample Output	<p>show ldp session brief on page 130</p> <p>show ldp session detail on page 130</p> <p>show ldp session extensive on page 130</p> <p>show ldp session auto-targeted detail on page 131</p>
Output Fields	Table 7 on page 126 describes the output fields for the show ldp session command. Output fields are listed in the approximate order in which they appear.

Table 7: show ldp session Output Fields

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

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

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

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

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

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

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

Sample Output

show ldp session brief

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

show ldp session detail

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

show ldp session extensive

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

Queue depth: 0

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

show ldp session auto-targeted detail

```

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

```

show ldp traffic-statistics


Syntax	<pre>show ldp traffic-statistics <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <p2mp></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>p2mp option added in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>
Description	Display Label Distribution Protocol (LDP) traffic statistics.
<div>  NOTE: If nonstop active routing features is configured, show ldp traffic-statistics command is not supported on backup Routing Engines. </div>	
Options	<p>none—Display LDP traffic statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display LDP traffic statistics for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>p2mp—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.</p>
Additional Information	To collect output from this command on a periodic basis, configure the traffic-statistics statement for the LDP protocol. For more information, see the <i>Junos MPLS Applications Configuration Guide</i> .
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ldp statistics Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs
List of Sample Output	<p>show ldp traffic-statistics on page 133</p> <p>show ldp traffic-statistics p2mp on page 134</p> <p>show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 134</p> <p>show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute) on page 134</p>
Output Fields	<p>Table 8 on page 133 lists the output fields for the show ldp traffic-statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 8: show ldp traffic-statistics Output Fields

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

Sample Output

show ldp traffic-statistics

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

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

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

show ldp traffic-statistics p2mp

```

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

```

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

```

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

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

```

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

```

user@host> show ldp traffic-statistics p2mp

```

P2MP FEC Statistics:

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

show ldp session

Syntax	<pre>show ldp session <brief detail extensive> <auto-targeted> <destination> <instance instance-name> <logical-system (all logical-system-name)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>auto-targeted option added in Junos OS Release 14.2.</p>
Description	Display information about Label Distribution Protocol (LDP) sessions.
Options	<p>none—Display standard information about all LDP sessions for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>auto-targeted—(Optional) Display information about LDP sessions that are automatically targeted using loopback addresses.</p> <p>destination—(Optional) Restrict LDP session display to the specified address.</p> <p>instance instance-name—(Optional) Display routing instance information for the specified instance. If instance-name is omitted, information is displayed for the master instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ldp session
List of Sample Output	<p>show ldp session brief on page 140</p> <p>show ldp session detail on page 140</p> <p>show ldp session extensive on page 140</p> <p>show ldp session auto-targeted detail on page 141</p>
Output Fields	Table 7 on page 126 describes the output fields for the show ldp session command. Output fields are listed in the approximate order in which they appear.

Table 9: show ldp session Output Fields

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

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

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

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

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

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

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

Sample Output

show ldp session brief

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

show ldp session detail

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

show ldp session extensive

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

Queue depth: 0

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

show ldp session auto-targeted detail

```

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

```

