Junos® OS

LDP Configuration Guide

Release
13.1

Published: 2013-02-28

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

Junos® OS LDP Configuration Guide
13.1

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

List of Figures

List of Tables

About the Documentation

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:

- T Series
- MX Series
- M Series
- PTX 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.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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.

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**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`

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

Convention	Description	Examples
Italic text like this	• Introduces or emphasizes important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub <default-metric** ***metric*****>;**
| (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** **(*string1* | *string2* | *string3*)**
# (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)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

CHAPTER 1

Introduction to LDP

LDP Introduction

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

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

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

Junos OS LDP Protocol Implementation

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

LDP Operation

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

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

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

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

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the Junos® OS Network Interfaces.

Tunneling LDP LSPs in RSVP LSPs

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

- Tunneling LDP LSPs in RSVP LSPs Overview on page 5
- Label Operations on page 5

Tunneling LDP LSPs in RSVP LSPs Overview

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

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

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

Label Operations

Figure 1 on page 6 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see Label Description.) 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

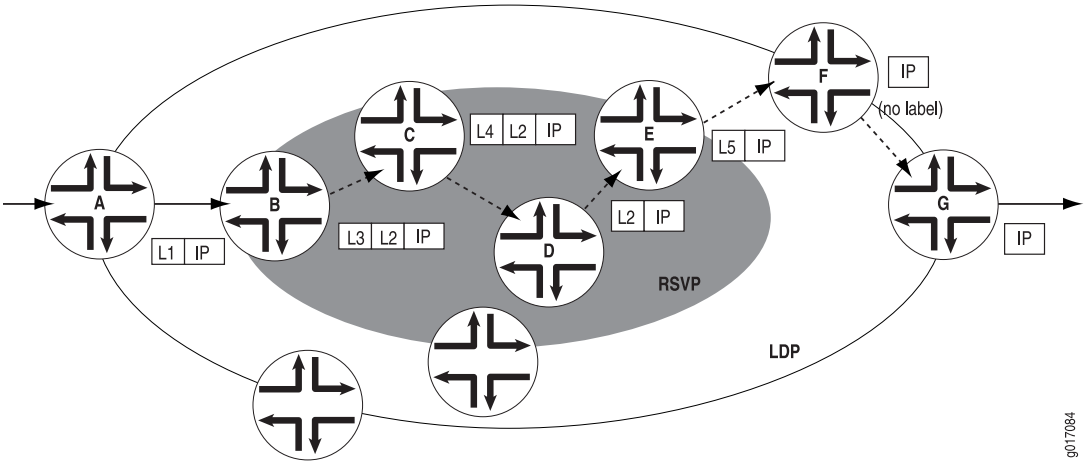

Figure 2 on page 6 depicts a double push label operation (L1L2). A double push label operation is used when the ingress router (A) for both the LDP LSP and the RSVP LSP tunneled through it is the same device. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

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

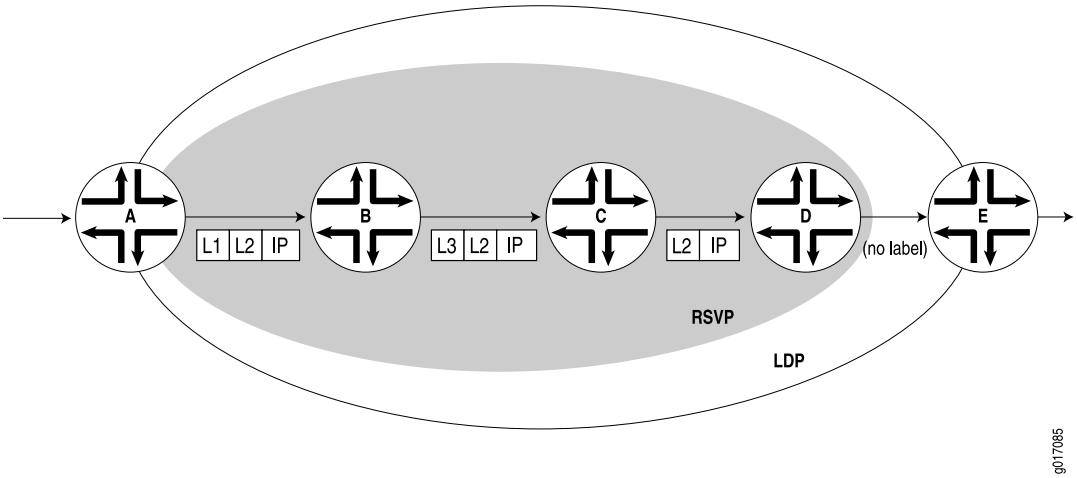

LDP Message Types

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

- Discovery Messages on page 7
- Session Messages on page 7
- Advertisement Messages on page 7
- Notification Messages on page 7

Discovery Messages

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

LDP uses the following discovery procedures:

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

Session Messages

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

Advertisement Messages

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

Notification Messages

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

- Error notifications, which signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.

- Advisory notifications, which pass information to a router about the LDP session or the status of some previous message received from the peer.

LDP Session Protection

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

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

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

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

LDP Graceful Restart

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

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

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an

initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

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

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

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.
- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

PART 2

Configuration

CHAPTER 2

LDP Configuration Guidelines

Minimum LDP Configuration

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

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

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

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

Enabling and Disabling LDP

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

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

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

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

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

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

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

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 15.

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 14.

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

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

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

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

Configuring the LDP Hold Time for Link Hello Messages

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

```
hold-time seconds;
```

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

Configuring the LDP Hold Time for Targeted Hello Messages

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

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

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

Enabling Strict Targeted Hello Messages for LDP

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

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

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

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

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

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

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

Configuring the Interval for LDP Keepalive Messages

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

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see "Configuring the Delay Before LDP Neighbors Are Considered Down" on page 15.

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

```
keepalive-interval seconds;
```

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

Configuring the LDP Keepalive Timeout

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

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

```
keepalive-timeout seconds;
```

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

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

Configuring LDP Route Preferences

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

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

```
preference preference;
```

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

Configuring LDP Graceful Restart

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

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

To configure LDP graceful restart, see the following sections:

- Enabling Graceful Restart on page 18
- Disabling LDP Graceful Restart or Helper Mode on page 19
- Configuring Reconnect Time on page 19
- Configuring Recovery Time and Maximum Recovery Time on page 20

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 Configuration Guide.

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

Disabling LDP Graceful Restart or Helper Mode

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

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

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

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

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

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

The following LDP graceful restart configurations are possible:

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

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

Configuring Reconnect Time

After the LDP connection between neighbors fails, neighbors wait a certain amount of time for the gracefully restarting router to resume sending LDP messages. After the wait

period, the LDP session can be reestablished. You can configure the wait period in seconds. This value is included in the fault tolerant session TLV sent in LDP initialization messages when LDP graceful restart is enabled.

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

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

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

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

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

Configuring Recovery Time and Maximum Recovery Time

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

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

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

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

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

Filtering Inbound LDP Label Bindings

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

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

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

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. Table 3 on page 21 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 Policy Configuration Guide.

Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

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

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

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

Filtering Outbound LDP Label Bindings

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

```
export [policy-name];
```

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

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in Table 4 on page 23.

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 Policy Configuration Guide.

Examples: Filtering Outbound LDP Label Bindings

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

```
[edit protocols]
ldp {
   export block-one;
```

```
}
policy-options {
   policy-statement block-one {
      term first {
         from {
            route-filter 10.10.255.6/32 exact;
         }
         then reject;
      }
      then accept;
   }
}
```

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

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

Specifying the Transport Address Used by LDP

You can control the transport address used by LDP. The transport address is the address used for the TCP session over which LDP is running. To configure transport address control, 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.

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 Configuration Guide.

Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

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

Configuring FEC Deaggregation

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

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

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

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

```
deaggregate;
```

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

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

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

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

```
no-deaggregate;
```

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

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

Related Documentation

- Configuring Load Balancing Across RSVP LSPs
- 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 Policy Configuration Guide.

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

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

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

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

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

Configuring LDP IPv4 FEC Filtering

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

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

If all the LDP neighbors associated with an LDP session are Layer 2 only, you can configure the Junos OS to advertise only Layer 2 circuit FECs by configuring the **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.

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;
    }
  }
  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 32.

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 Junos OS Operational Mode Commands.

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.

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 32.) 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-livenesss-detection]** hierarchy level or at the **[edit protocols ldp oam fec address bfd-livenesss-detection]** hierarchy level. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

```
holddown-interval seconds;
```

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

Configuring 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**]**

Configuring LDP LSP Traceroute

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

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

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

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

You can configure this statement at the following hierarchy levels:

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

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

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

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 34
- Disabling LDP Statistics on the Penultimate-Hop Router on page 35
- LDP Statistics Limitations on page 35

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

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;
        }
    }
}
```

Configuring Miscellaneous LDP Properties

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

Configuring LDP to Use the IGP Route Metric

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

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

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

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

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

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

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

```
no-forwarding;
```

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

Multiple-Instance LDP and Carrier-of-Carriers VPNs

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

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

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

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

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

```
explicit-null;
```

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

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

For more information about labels, see Label Description and Label Allocation.

Enabling LDP over RSVP-Established LSPs

You can run LDP over LSPs established by RSVP, effectively tunneling the LDP-established LSP through the one established by RSVP. To do so, enable LDP on the lo0.0 interface (see "Enabling and Disabling LDP" on page 14). You must also configure the LSPs over which you want LDP to operate by including the **ldp-tunneling** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

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

```
}
```

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

Related Documentation

- Tunneling LDP LSPs in RSVP LSPs Overview on page 5

Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

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

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

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

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

You can configure this statement at the following hierarchy levels:

- **[edit protocols ospf traffic-engineering shortcuts]**
- **[edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]**

To enable LDP over RSVP LSPs, you also still need to complete the procedure in Section "Enabling LDP over RSVP-Established LSPs" on page 40.

Configuring the TCP MD5 Signature for LDP Sessions

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Configuring LDP Session Protection

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

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

```
session-protection {
```

```
    timeout seconds;
}
```

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

Disabling SNMP Traps for LDP

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

For information about the LDP SNMP traps and the proprietary LDP MIB, see the Network Management Configuration Guide.

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

```
log-updown {
    trap disable;
}
```

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

Configuring LDP Synchronization with the IGP on LDP Links

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

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

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

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

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

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

Configuring LDP Synchronization with the IGP on the Router

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

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

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

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

Configuring the Label Withdrawal Timer

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

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

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

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

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

Ignoring the LDP Subnet Check

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

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

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

This statement can be included at the following hierarchy levels:

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

Configuring Multicast LDP Link Protection

A Layer Distribution Protocol (LDP) label-switched path (LSP) that is point to multipoint can be used to send traffic from a single root or ingress node to a number of leaf or egress nodes traversing one or more transit nodes. Multicast Label Distribution Protocol (MLDP) link protection enables fast reroute of traffic carried over point-to-multipoint LDP LSPs in case of a link failure. When one of the links of the point-to-multipoint tree fails, the subtrees might get detached until the IGP reconverges and MLDP initiates label mapping using the best path from the downstream router to the new upstream router.

To protect the traffic flowing through the LDP point-to-multipoint LSP, you can configure an explicit tunnel for traffic to be re-routed in the event of link failure. The explicit path has to terminate on the next downstream router, and the reverse path forwarding for the traffic should be successful.

You can configure MLDP link protection using dynamic RSVP LSPs or a regular LDP of unidirectional paths using hop-by-hop routing. Dynamic RSVP LSPs are used as a bypass tunnel. The RSVP LSP's explicit route object (ERO) is calculated using Constrained Shortest Path First (CSPF) with the constraint as the link to avoid. The LSP is signaled and torn down dynamically whenever link protection is necessary. A targeted adjacency to the downstream label-switching router (LSR) is created (if none is preconfigured already) for two reasons:

- Keeping the session up after link failure.
- Using the point-to-multipoint label received during the session to send traffic to the downstream LSR using RSVP LSP as a bypass tunnel.

It is possible to configure a remote neighbor to the LSR by configuring LDP tunneling on RSVP LSPs, or LDP-based virtual private LAN service (VPLS), or on Layer 2 circuits, or LDP session protection.

To enable MLDP link protection, Junos OS supports the make-before-break (MBB) feature to ensure minimum packet loss when attempting to signal a new LSP path before tearing down the old LSP path.

An LSR selects its upstream LSR as its next hop to the root of a point-to-multipoint LSP. When the best path to reach the root changes, the LSR chooses a new upstream LSR. During this transition, the LSP might go down temporarily resulting in packet loss until the LSP reconverges to the new upstream LSR. By configuring MBB, you can minimize packet loss during reconvergence. In addition, there might be scenarios where the best path from LSR to the root changes and yet the LSP continues to forward packets to the earlier next hop to the root. In such cases, a new LSP must be established before the old LSP is brought down to minimize the duration of packet loss. If a link fails, the downstream LSR continues to receive and forward packets to other downstream LSRs as it continues to receive packets from the RSVP LSP.

NOTE: You must configure link protection for the LDP interface using the **link-protection** statement at the **[edit protocols ldp]** hierarchy level before configuring MBB.

To configure make before break, include the **make-before-break** statement at the **[edit protocols ldp]** hierarchy level:

```
make-before-break {
    timeout seconds;
    switchover-delay seconds;
}
```

When you include the **make-before-statement** in the configuration, the LDP LSR advertises that it is capable of handling MBB point-to-multipoint LSPs configured using the **p2mp** configuration statement at the **[edit protocols ldp]** hierarchy level.

You can include the following options for the **make-before-break** statement:

- **switchover-delay**—Specify a value from 1 through 300 seconds to change switchover delay for a point-to-multipoint LSP from the old LSR to the new upstream LSR. The default value is 30 seconds. If an MBB acknowledgement 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.
- **timeout**—Specify a value from 1 through 300 seconds to change make-before-break timeout for point-to-multipoint LSPs. The default value is 30 seconds. Even if an MBB acknowledgment is not received for a point-to-multipoint LSP before the specified timeout period expires, the LSR performs an MBB switchover from the old LSR to the new upstream LSR.

Related Documentation

- make-before-break (LDP) on page 81

CHAPTER 3

LDP Example

Example: Configuring LDP Downstream on Demand

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

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

Requirements

This example uses the following hardware and software components:

- M Series router
- Junos OS 12.2

Overview

You can enable LDP downstream on demand label advertisement for an LDP session by including the **downstream-on-demand** statement at the **[edit protocols ldp session]** hierarchy level. If you have configured downstream on demand, the Juniper Networks router advertises the downstream on demand request to its peer routers. For a downstream on demand session to be established between two routers, both have to

advertise downstream on demand mode during LDP session establishment. If one router advertises downstream unsolicited mode and the other advertises downstream on demand, downstream unsolicited mode is used.

Configuration

Configuring LDP Downstream on Demand

Step-by-Step Procedure

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

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

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

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

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

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

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

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

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

Distributing LDP Downstream on Demand Routes into Labeled BGP

Step-by-Step Procedure

To distribute LDP downstream on demand routes into labeled BGP, use a BGP export policy.

1. Configure the LDP route policy (**redistribute_ldp** in this example).

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

```
user@host# set policy-statement redistribute_ldp term 1 then accept
```

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

   BGP forwards the LDP routes based on the **redistribute_ldp** policy to the remote PE router

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

Step-by-Step Procedure

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

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

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

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

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

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

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

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

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

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

```
user@host#

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

user@host#

show protocols ldp
dod-request-policy DOD-Request-Loopbacks;
session 1.1.1.1 {
    downstream-on-demand;
}

user@host#

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


Verification

Verifying Label Advertisement Mode

Purpose Confirm that the configuration is working properly.

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

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

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

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

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

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

PART 3

Administration

CHAPTER 4

LDP Standards

Supported LDP Standards

Junos OS substantially supports the following RFCs, which define standards for LDP.

- RFC 3212, *Constraint-Based LSP Setup using LDP*
- RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol*

The following RFCs do not define standards, but provide information about LDP. The IETF classifies them as "Informational."

- RFC 3215, *LDP State Machine*
- RFC 5036, *LDP Specification*

  For the following features described in the indicated sections of the RFC, Junos OS supports one of the possible modes but not the others:

  - Label distribution control (section 2.6.1): Ordered mode is supported, but not Independent mode.
  - Label retention (section 2.6.2): Liberal mode is supported, but not Conservative mode.
  - Label advertisement (section 2.6.3): Downstream Unsolicited mode is supported, but not Downstream on Demand mode.
- RFC 5443, *LDP IGP Synchronization*

Related Documentation

- Supported GMPLS Standards
- Supported MPLS Standards
- Supported RSVP Standards
- Accessing Standards Documents on the Internet

CHAPTER 5

Summary of LDP Configuration Statements

allow-subnet-mismatch

Syntax allow-subnet-mismatch;

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

Release Information Statement introduced in Junos OS Release 9.3.

Description Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.

Default The source address in the LDP link hello packet is matched against the interface address.

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

Related Documentation

- Ignoring the LDP Subnet Check on page 44

authentication-algorithm

Syntax authentication-algorithm *algorithm*;

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

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced for BGP in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure an authentication algorithm type.

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.

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

Related Documentation
- Understanding Route Authentication
- Example: Configuring Route Authentication for BGP

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.

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

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

Related Documentation
- Configuring the TCP MD5 Signature for LDP Sessions on page 41

authentication-key-chain (Protocols LDP)

Syntax authentication-key-chain *key-chain*;

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

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the **0.0.0.0/allow** statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.

Options *key-chain*—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

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

Related Documentation
- Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols
- Configuring Miscellaneous LDP Properties on page 39

bfd-liveness-detection (Protocols LDP)

Syntax

```
bfd-liveness-detection {
  detection-time threshold milliseconds;
  ecmp;
  failure-action {
    remove-nexthop;
    remove-route;
  }
  holddown-interval seconds;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  minimum-transmit-interval milliseconds;
  multiplier detection-time-multiplier;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
}
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 7.6.
Support for the **bfd-liveness-detection** statement at the **[edit protocols ldp oam fec *address*]** hierarchy level and the **ecmp** option added in Junos OS Release 9.0.
Support for the **failure-action** statement with the **remove-nexthop** and **remove-route** options and the **holddown-interval** statement added in Junos OS Release 9.4.

Description

Enable Bidirectional Forwarding Detection (BFD) for all MPLS LSPs or for just a specific LSP.

Options

minimum-interval—Minimum transmit and receive interval.

Range: 50 through 255,000 milliseconds

Default: 50

minimum-receive-interval—Minimum receive interval.

Range: 50 through 255,000 milliseconds

Default: 50

minimum-transmit-interval—Minimum transmit interval.

Range: 50 through 255,000 milliseconds

Default: 50

multiplier—Detection time multiplier.

Range: 50 through 255

Default: 3

The other options are explained separately.

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

Related Documentation
- Configuring BFD for LDP LSPs on page 28

deaggregate

Syntax
deaggregate | no-deaggregate;

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

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Control forwarding equivalence class (FEC) deaggregation on the router. The use of the **deaggregate** statement in LDP is a standard practice that we recommend for LDP deployments.

Default
Deaggregation is disabled on the router.

Options
deaggregate—Deaggregate FECs.

no-deaggregate—Aggregate FECs.

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

Related Documentation
- Configuring FEC Deaggregation on page 26

disable (Protocols LDP)

Syntax disable;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.

Default LDP is enabled on interfaces configured with the LDP **interface** statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the **[edit routing-options]** hierarchy level.

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

Related Documentation
- Enabling and Disabling LDP on page 14
- Configuring LDP Graceful Restart on page 18

dod-request-policy

Syntax dod-request-policy *dod-request-policy-name*;

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

Release Information Statement introduced in Junos OS Release 12.2.

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 ***dod-request-polcy-name***—Specify the name of the downstream on demand request policy.

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

Related Documentation
- Example: Configuring LDP Downstream on Demand on page 49

downstream-on-demand

Syntax downstream-on-demand;

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

Release Information Statement introduced in Junos OS Release 12.2.

Description Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.

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

Related Documentation
- Example: Configuring LDP Downstream on Demand on page 49

ecmp

Syntax ecmp;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp oam bfd-liveness-detection],
[edit logical-systems *logical-system-name* protocols ldp oam fec *address* bfd-liveness-detection],
[edit protocols ldp oam bfd-liveness-detection],
[edit protocols ldp oam fec *address* bfd-liveness-detection]

Release Information Statement introduced in Junos OS Release 8.5.

Description Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the **ecmp** statement, you must also configure the **periodic-traceroute** statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the **periodic-traceroute** statement at the global hierarchy level (**[edit protocols ldp oam]**) while only configuring the **ecmp** statement for a specific FEC (**[edit protocols ldp oam fec *address* bfd-liveness-detection]**).

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

Related Documentation
- Configuring ECMP-Aware BFD for LDP LSPs on page 31

egress-policy

Syntax egress-policy [*policy-names*];

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Control the prefixes advertised into LDP.

Default Only the loopback address is advertised.

Options ***policy-names***—Name of one or more routing policies.

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

Related Documentation
- Configuring the Prefixes Advertised into LDP from the Routing Table on page 25

explicit-null (Protocols LDP)

Syntax explicit-null;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Advertise label 0 to the egress router of a label-switched path (LSP).

Default If you do not include the **explicit-null** statement in the MPLS configuration, label 3 (implicit null) is advertised.

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

Related Documentation
- Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 40

export (Protocols LDP)

Syntax export [*policy-names*];

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options ***policy-names***—Name of one or more routing policies.

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

Related Documentation
- Filtering Outbound LDP Label Bindings on page 22

failure-action (Protocols LDP)

Syntax

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

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ldp oam bfd-livenesss-detection],
[edit logical-systems *logical-system-name* protocols ldp oam fec *address* bfd-livenesss-detection],
[edit protocols ldp oam bfd-livenesss-detection],
[edit protocols ldp oam fec *address* bfd-livenesss-detection]

Release Information

Statement introduced in Junos OS Release 9.4.

Description

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

Options

remove-nexthop—Remove a route corresponding to a next hop of the LSP's route at the ingress node when a BFD session failure event is detected.

remove-route—Remove the route corresponding to an LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

Required Privilege Level

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

Related Documentation

- Configuring a Failure Action for the BFD Session on an LDP LSP on page 31

fec

Syntax

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

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

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).

Options ***fec-address***—Specify the FEC address.

The other statements are explained separately.

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

Related Documentation

- Configuring BFD for LDP LSPs on page 28

graceful-restart (Protocols LDP)

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  maximum-neighbor-recovery-time value;
  reconnect-time seconds;
  recovery-time value;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable 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 18

hello-interval (Protocols LDP)

Syntax hello-interval *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Support for LDP targeted hellos added in Junos OS Release 9.5.

Description Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the **hello-interval** statement.

Options ***seconds***—Length of time between transmission of hello packets.

Range: 1 through 65,535 seconds

Default: 5 seconds for link hello messages, 15 seconds for targeted hello messages

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

Related Documentation
- Configuring the LDP Timer for Hello Messages on page 14

helper-disable (LDP)

Syntax helper-disable;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.

Default Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.

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

Related Documentation
- Configuring LDP Graceful Restart on page 18

holddown-interval

Syntax holddown-interval *holddown-interval*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp oam bfd-livenesss-detection],
[edit logical-systems *logical-system-name* protocols ldp oam fec *address* bfd-livenesss-detection],
[edit protocols ldp oam bfd-livenesss-detection],
[edit protocols ldp oam fec *address* bfd-livenesss-detection]

Release Information Statement introduced in Junos OS Release 9.4.

Description Specify how long the BFD session should be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

Options ***holddown-interval***—Number of seconds the BFD session should remain up before adding the route or next hop.

Default: 0 seconds

Range: 0 through 65,535 seconds

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

Related Documentation
- Configuring the Holddown Interval for the BFD Session on page 32

hold-time (Protocols LDP)

Syntax hold-time *seconds*;

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

Release Information
Statement introduced before Junos OS Release 7.4.
Support for LDP targeted hellos added in Junos OS Release 9.5.

Description Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the **hold-time** statement.

Options ***seconds***—Hold-time value.

Range: 1 through 65,535 seconds

Default: 15 seconds for link hello messages, 45 seconds for targeted hello messages

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

Related Documentation
- Configuring the Delay Before LDP Neighbors Are Considered Down on page 15

ignore-lsp-metrics

Syntax ignore-lsp-metrics;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts],
[edit protocols ospf traffic-engineering shortcuts]

Release Information Statement introduced in Junos OS Release 7.5.

Description Cause OSPF to ignore the RSVP LSP metric.

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.

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

Related Documentation
- Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 41

igp-synchronization

Syntax igp-synchronization holddown-interval *seconds*;

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

Release Information Statement introduced in Junos OS Release 9.5.

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

Options **holddown-interval *seconds***—Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational.

Default: 10 seconds

Range: 10 through 60 seconds

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

Related Documentation
- Configuring LDP Synchronization with the IGP on the Router on page 44

import (Protocols LDP)

Syntax import [*policy-names*];

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.

Options ***policy-names***—Name of one or more routing policies.

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

Related Documentation
- Filtering Inbound LDP Label Bindings on page 20

ingress-policy

Syntax ingress-policy [*ingress-policy-names*];

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

Release Information Statement introduced in Junos OS Release 9.4.

Description 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.

Options ***ingress-policy-names***—Specify the names of the ingress policies.

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

Related Documentation
- Configuring OAM Ingress Policies for LDP on page 32

interface (Protocols LDP)

Syntax

```
interface interface-name {
    disable;
    hello-interval seconds;
    hold-time seconds;
    transport-address (interface | loopback);
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable LDP on one or more router interfaces.

Default

LDP is disabled on all interfaces.

Options

interface-name—Name of an interface. To configure all interfaces, specify **all**.

The remaining statements are explained separately.

Required Privilege Level

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

Related Documentation

- Enabling and Disabling LDP on page 14

keepalive-interval

Syntax keepalive-interval *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the keepalive interval value.

Options ***seconds***—Keepalive value.

Range: 1 through 65,535

Default: 10 seconds

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

Related Documentation
- Configuring the Interval for LDP Keepalive Messages on page 17

keepalive-timeout

Syntax keepalive-timeout *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.

Options ***seconds***—Keepalive timeout value.

Range: 1 through 65,535

Default: 30 seconds

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

Related Documentation
- Configuring the LDP Keepalive Timeout on page 17

l2-smart-policy

Syntax l2-smart-policy;

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

Release Information Statement introduced in Junos OS Release 8.4.

Description Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.

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

Related Documentation
- Configuring LDP IPv4 FEC Filtering on page 27

label-withdrawal-delay

Syntax label-withdrawal-delay *seconds*;

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

Release Information Statement introduced in Junos OS Release 9.1.

Description Delay the withdrawal of labels to reduce router workload during IGP convergence.

Options ***seconds***—Configure the number of seconds to wait before withdrawing labels for the LDP LSPs.

Default: 60 seconds

Range: 0 through 300 seconds

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

Related Documentation
- Configuring the Label Withdrawal Timer on page 44

ldp

Syntax ldp { ... }

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.

Description Enable LDP routing on the router or switch.

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

Default LDP is disabled on the router.

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

Related Documentation

- Minimum LDP Configuration on page 14
- Enabling and Disabling LDP on page 14

ldp-synchronization

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 7.5.

Description

Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.

Options

The other statements are explained separately.

Required Privilege Level

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

Related Documentation

- Configuring LDP Synchronization with the IGP on LDP Links on page 43

log-updown (Protocols LDP)

Syntax

```
log-updown {
    trap disable;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Disable LDP traps on the router, logical system, or routing instance.

Options

trap disable—Disable LDP traps.

Default: LDP traps are enabled on the router.

Required Privilege Level

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

Related Documentation

- Disabling SNMP Traps for LDP on page 43

make-before-break (LDP)

Syntax

```
make-before-break {
  timeout seconds;
  switchover-delay seconds;
}
```

Hierarchy Level [edit protocols ldp]

Release Information Statement introduced in Junos OS Release 12.3.

Description Configures make before break (MBB) for multicast LDP (MLDP) link protection to ensure minimum packet loss when attempting to signal a new label-switched path (LSP) before tearing down the old LSP path.

Options

timeout *seconds*—Specify a value to change a make -before-break timeout for point-to-multipoint LSPs. Even if an MBB acknowledgment is not received for a point-to-multipoint LSP before the specified timeout period expires, the label-switching router (LSR) performs an MBB switchover from the old LSR to the new upstream LSR.

Range: 1 through 300 seconds

Default: 30 seconds

switchover-delay *seconds*—Specify a value to change switchover delay for a point-to-multipoint LSP from the old LSR to the new upstream LSR. If an MBB acknowledgment is received on a point of local repair (PLR) router, the PLR waits for the specified seconds to switch its upstream LSR from the old LSR to the new LSR.

Range: 1 through 300 seconds

Default: 30 seconds

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

Related Documentation

- Configuring Multicast LDP Link Protection on page 46

maximum-neighbor-recovery-time

Syntax maximum-neighbor-recovery-time *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4. Statement changed from **maximum-recovery-time** to **maximum-neighbor-recovery-time** in Junos OS Release 9.1.

Description Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.

Options ***seconds***—Configure the maximum recovery time, in seconds.

Range: 120 through 1800 seconds

Default: 140 seconds

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

Related Documentation

- Configuring Recovery Time and Maximum Recovery Time on page 20
- Configuring Graceful Restart Options for LDP
- no-strict-lsa-checking
- recovery-time

no-forwarding

Syntax no-forwarding;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Do not add ingress routes to the inet.0 routing table even if **traffic-engineering bgp-igp** (configured at the **[edit protocols mpls]** hierarchy level) is enabled.

Default The **no-forwarding** statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when **traffic-engineering bgp-igp** is enabled.

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

Related Documentation
- Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 39
- Configuring Virtual-Router Routing Instances in VPNs

oam (Protocols LDP)

Syntax

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

Hierarchy Level

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

Release Information

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

Description

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

Options

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

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

Default: 60 seconds

Range: 30 through 3,600 seconds

The remaining statements are explained separately.

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 28

p2mp (Protocols LDP)

Syntax
p2mp;

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

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 8.4.
Support added at the **[edit protocols ldp oam]** and **[edit logical-systems *logical-system-name* protocols ldp oam]** hierarchy levels in Junos OS Release 9.0.
Statement introduced in Junos OS Release 12.2 for EX Series switches.

Description

Enable tracing of forwarding equivalence classes (FECs) for LDP LSPs.

Options

disable—(Optional) Disable tracing for a specific FEC. This option is available at the **[edit protocols ldp oam fec *fec-address* periodic-traceroute]** and **[edit logical-systems *logical-system-name* protocols ldp oam fec *fec-address* periodic-traceroute]** hierarchy levels only.

exp *exp-value*—(Optional) Specify the class of service to use when sending probes.

Default: 7

Range: 0 through 7

fanout *fanout-value*—(Optional) Specify the maximum number of next hops to search per node.

Default: 16

Range: 1 though 16

frequency *minutes*—(Optional) Specify the interval between traceroute attempts.

Default: 60 minutes

Range: 15 through 120 minutes

paths *number-of-paths*—(Optional) Specify the maximum number of paths to search.

Default: 3

Range: 1 through 255

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

Default: 3

Range: 1 through 9

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

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

Default: 64

Range: 1 through 255

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

Default: 10 seconds

Range: 5 though 15 seconds

Required Privilege Level

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

Related Documentation

- Configuring LDP LSP Traceroute on page 32

policing (Protocols LDP)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable policing of forwarding equivalence classes (FECs) for LDP.

Options

fec *fec-address*—Specify the address for the FEC.

ingress-traffic *filter-name*—Specify the name of the filter for policing ingress FEC traffic.

transit-traffic *filter-name*—Specify the name of the filter for policing transit FEC traffic.

Required Privilege Level

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

Related Documentation

- Configuring Policers for LDP FECs on page 26

preference (Protocols LDP)

Syntax preference *preference*;

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 interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ldp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the route preference level for LDP routes.

Options ***preference***—Preferred value.
Range: 0 through 255
Default: 9

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

Related Documentation
- Configuring LDP Route Preferences on page 18

reconnect-time

Syntax reconnect-time *seconds*;

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

Release Information Statement introduced in Junos OS Release 9.1.

Description Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.

Options ***seconds***—Time required for reconnection.
Range: 30 through 300
Default: 60 seconds

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

Related Documentation
- Configuring LDP Graceful Restart on page 18 on LDP Configuration Guide
- Configuring Graceful Restart Options for LDP

recovery-time

Syntax recovery-time *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the amount of time a router waits for LDP to restart gracefully.

Options ***seconds***—Configure the recovery time, in seconds.

Range: 120 through 1800 seconds

Default: 140 seconds

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

Related Documentation
- Configuring Recovery Time and Maximum Recovery Time on page 20

session (ldp)

Syntax
```
session address {
    authentication-algorithm algorithm;
    authentication-key authentication-key;
    authentication-key-chain key-chain-name;
}
```

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

Release Information
Statement introduced before Junos OS Release 7.4.
authentication-algorithm statement introduced in Junos OS Release 7.6.

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
- Configuring the TCP MD5 Signature for LDP Sessions on page 41

session-protection

Syntax

```
session-protection {
    timeout seconds;
}
```

Hierarchy Level

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

Description

Configure when an LDP session is torn down and resignaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.

Options

timeout *seconds*—Time in seconds before the LDP session is torn down and resignaled.

Range: 1 through 65,535 seconds

Required Privilege Level

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

Related Documentation

- Configuring LDP Session Protection on page 42

strict-targeted-hellos

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.

Required Privilege Level

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

Related Documentation

- Enabling Strict Targeted Hello Messages for LDP on page 17

targeted-hello

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 9.5.

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

- Configuring the LDP Timer for Hello Messages on page 14
- Configuring the Delay Before LDP Neighbors Are Considered Down on page 15

traceoptions (Protocols LDP)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
match-on address option for the **filter** flag modifier added in Junos OS Release 10.4.

Description

LDP protocol-level trace options.

Default

The default LDP protocol-level trace options are inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **ldp-log**. We recommend that you place LDP tracing output in the file **ldp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000

Default: 2 files

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

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **address**—Operation of address and address withdrawal messages
- **binding**—Label-binding operations
- **error**—Error conditions
- **event**—Protocol events
- **initialization**—Operation of initialization messages

- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **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: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

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

Default: 1 MB

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

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

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

Related Documentation
- Tracing LDP Protocol Traffic on page 36
- Network Management Configuration Guide

track-igp-metric

Syntax
track-igp-metric;

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

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

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

Related Documentation
- Configuring LDP to Use the IGP Route Metric on page 39

traffic-statistics (Protocols LDP)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.

Options

file *filename*—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of LDP statistics files. When a statistics file named ***ldp-stat*** reaches its maximum size, it is renamed ***ldp-stat*.0**, then ***ldp-stat*.1**, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.

Range: 2 through 1000

Default: 2 files

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

interval *seconds*—(Optional) Specify the interval at which the statistics are polled and written to the file.

Default: 300 seconds (5 minutes)

no-penultimate-hop—(Optional) Do not collect traffic statistics on the penultimate hop router.

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

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

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

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

Default: 1 MB

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

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

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

Related Documentation
- Collecting LDP Statistics on page 33

transport-address

Syntax
transport-address (interface | router-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 interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ldp interface *interface-name*]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Enable control of the transport address used by LDP.

Default
router-id

Options
interface—The first IP address on the interface is used as the transport address.

router-id—The router identifier is used as the transport address.

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

Related Documentation
- Specifying the Transport Address Used by LDP on page 24

PART 4

Index

Index

S

T

W