



---

# System Services Feature Guide for the QFX Series

Release  
15.1



---

Modified: 2016-03-02

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

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

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

*System Services Feature Guide for the QFX Series*

15.1

Copyright © 2016, Juniper Networks, Inc.

All rights reserved.

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

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

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

# Table of Contents

	About the Documentation . . . . .	xi
	Documentation and Release Notes . . . . .	xi
	Supported Platforms . . . . .	xi
	Using the Examples in This Manual . . . . .	xi
	Merging a Full Example . . . . .	xii
	Merging a Snippet . . . . .	xii
	Documentation Conventions . . . . .	xiii
	Documentation Feedback . . . . .	xv
	Requesting Technical Support . . . . .	xv
	Self-Help Online Tools and Resources . . . . .	xv
	Opening a Case with JTAC . . . . .	xvi
<b>Part 1</b>	<b>Port Mirroring</b>	
<b>Chapter 1</b>	<b>Configuring Port Mirroring . . . . .</b>	<b>3</b>
	Understanding Port Mirroring . . . . .	3
	Port Mirroring Overview . . . . .	3
	Port Mirroring Instance Types . . . . .	4
	Port-Mirroring Terminology . . . . .	4
	Port Mirroring and STP . . . . .	6
	Port Mirroring Constraints and Limitations . . . . .	6
	Local and Remote Port Mirroring . . . . .	6
	Remote Port Mirroring Only . . . . .	8
	Configuring Port Mirroring . . . . .	8
	Configuring Port Mirroring for Local Analysis . . . . .	9
	Configuring Port Mirroring for Remote Analysis . . . . .	9
	Filtering the Traffic Entering an Analyzer . . . . .	10
	Examples: Configuring Port Mirroring for Local Analysis . . . . .	11
	Example: Mirroring Employee Web Traffic with a Firewall Filter . . . . .	13
	Example: Configuring Port Mirroring for Remote Analysis . . . . .	16
	Example: Mirroring Employee Web Traffic with a Firewall Filter . . . . .	21
	Troubleshooting Port Mirroring . . . . .	24
	Port Mirroring Constraints and Limitations . . . . .	24
	Local and Remote Port Mirroring . . . . .	24
	Remote Port Mirroring Only . . . . .	26
	Egress Port Mirroring with VLAN Translation . . . . .	26
	Egress Port Mirroring with Private VLANs . . . . .	26

<b>Part 2</b>	<b>DHCP and DHCP Relay</b>	
<b>Chapter 2</b>	<b>Using DHCP and DHCP Relay</b>	<b>31</b>
	DHCP and BOOTP Relay Overview	31
	Configuring DHCP and BOOTP	31
	Verifying and Managing DHCP Local Server Configuration	32
	Verifying and Managing DHCPv6 Local Server Configuration	33
	Verifying and Managing DHCP Relay Configuration	33
	Verifying and Managing DHCPv6 Relay Configuration	33
	Understanding Layer 3 Logical Interfaces	34
<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 3</b>	<b>Configuration Statements (Port Mirroring)</b>	<b>37</b>
	analyzer	38
	egress	39
	ethernet-switching (Port Mirroring)	40
	family (Port Mirroring)	41
	inet (Port Mirroring)	42
	ingress (Port Mirroring)	43
	input	44
	instance (Port Mirroring)	45
	interface (Port Mirroring)	46
	ip-address (Port Mirroring)	47
	output	48
	port-mirroring	49
	routing-instance (Port Mirroring)	50
	vlan (Port Mirroring)	51
<b>Chapter 4</b>	<b>Configuration Statements (DHCP and DHCP Relay)</b>	<b>53</b>
	dhcp-local-server	54
	dhcp-relay	60
	dhcp-service	65
	forwarding-options	66
	forward-only (DHCP Relay Agent)	71
<b>Chapter 5</b>	<b>Configuration Statements (Encryption)</b>	<b>73</b>
	authentication-key-chains	74
	cache-size	75
	cache-timeout-negative	76
	ca-name	77
	certificates	78
	certification-authority	79
	crl (Encryption Interface)	80
	encoding	80
	enrollment-retry	81
	enrollment-url	81
	file	82
	key (Authentication Keychain)	83
	key-chain (Security)	84

---

	ldap-url . . . . .	85
	local . . . . .	86
	maximum-certificates . . . . .	87
	path-length . . . . .	87
	secret . . . . .	88
	security . . . . .	89
	ssh-known-hosts . . . . .	90
	start-time (Authentication Key Transmission) . . . . .	91
	traceoptions . . . . .	92
<b>Chapter 6</b>	<b>Operational Command (Port Mirroring) . . . . .</b>	<b>95</b>
	show analyzer . . . . .	96
<b>Chapter 7</b>	<b>Operational Commands (DHCP Local Server) . . . . .</b>	<b>99</b>
	clear dhcp server binding . . . . .	100
	clear dhcp server statistics . . . . .	103
	clear dhcpv6 server binding . . . . .	105
	clear dhcpv6 server statistics . . . . .	107
	request dhcp server reconfigure . . . . .	108
	request dhcpv6 server reconfigure . . . . .	110
	request system reboot . . . . .	112
	show dhcp server binding . . . . .	117
	show dhcp server statistics . . . . .	123
	show dhcpv6 server binding . . . . .	126
	show dhcpv6 server statistics . . . . .	132
<b>Chapter 8</b>	<b>Operational Commands (DHCP Relay Agent) . . . . .</b>	<b>135</b>
	clear dhcp relay binding . . . . .	136
	clear dhcp relay statistics . . . . .	138
	clear dhcpv6 relay binding . . . . .	141
	clear dhcpv6 relay statistics . . . . .	144
	show dhcp relay binding . . . . .	146
	show dhcp relay statistics . . . . .	151
	show dhcpv6 relay binding . . . . .	154
	show dhcpv6 relay statistics . . . . .	160
	show route extensive . . . . .	163
	show route protocol . . . . .	180



# List of Figures

Part 1	Port Mirroring	
Chapter 1	Configuring Port Mirroring .....	3
	Figure 1: Network Topology for Local Port Mirroring Example .....	12





# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xi</b>
	Table 1: Notice Icons . . . . .	xiii
	Table 2: Text and Syntax Conventions . . . . .	xiii
<b>Part 1</b>	<b>Port Mirroring</b>	
<b>Chapter 1</b>	<b>Configuring Port Mirroring</b> . . . . .	<b>3</b>
	Table 3: Port Mirroring Terms and Definitions . . . . .	4
<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 6</b>	<b>Operational Command (Port Mirroring)</b> . . . . .	<b>95</b>
	Table 4: show analyzer Output Fields . . . . .	96
<b>Chapter 7</b>	<b>Operational Commands (DHCP Local Server)</b> . . . . .	<b>99</b>
	Table 5: show dhcp server binding Output Fields . . . . .	118
	Table 6: show dhcp server statistics Output Fields . . . . .	124
	Table 7: show dhcpv6 server binding Output Fields . . . . .	127
	Table 8: show dhcpv6 server statistics Output Fields . . . . .	133
<b>Chapter 8</b>	<b>Operational Commands (DHCP Relay Agent)</b> . . . . .	<b>135</b>
	Table 9: clear dhcp relay statistics Output Fields . . . . .	139
	Table 10: show dhcp relay binding Output Fields . . . . .	147
	Table 11: show dhcp relay statistics Output Fields . . . . .	152
	Table 12: show dhcpv6 relay binding Output Fields . . . . .	155
	Table 13: show dhcpv6 relay statistics Output Fields . . . . .	160
	Table 14: show route extensive Output Fields . . . . .	163



# About the Documentation

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

## Documentation and Release Notes

---

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

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

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

## Supported Platforms

---

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

- QFX Series standalone switches

## Using the Examples in This Manual

---

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

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

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

## Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

## Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

## Documentation Conventions

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

Table 1: Notice Icons

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

Table 2 on page xiii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"><li>Introduces or emphasizes important new terms.</li><li>Identifies guide names.</li><li>Identifies RFC and Internet draft titles.</li></ul>	<ul style="list-style-type: none"><li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li><li><i>Junos OS CLI User Guide</i></li><li>RFC 1997, <i>BGP Communities Attribute</i></li></ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"><li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [</b> <i>community-ids</i> <b>]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>

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

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

## Documentation Feedback

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

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

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

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

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

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

## Opening a Case with JTAC

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

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

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



## PART 1

# Port Mirroring

- [Configuring Port Mirroring on page 3](#)



## CHAPTER 1

# Configuring Port Mirroring

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 16](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 21](#)
- [Troubleshooting Port Mirroring on page 24](#)

## Understanding Port Mirroring

---

- [Port Mirroring Overview on page 3](#)
- [Port Mirroring Instance Types on page 4](#)
- [Port-Mirroring Terminology on page 4](#)
- [Port Mirroring and STP on page 6](#)
- [Port Mirroring Constraints and Limitations on page 6](#)

### Port Mirroring Overview

Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring or to a VLAN for remote monitoring. Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

Port mirroring is needed for traffic analysis on a switch because a switch normally sends packets only to the port to which the destination device is connected. You configure port mirroring on the switch to send copies of unicast traffic to a local interface or a VLAN and run an analyzer application on a device connected to the interface or VLAN. You configure port mirroring by using the **analyzer** statement.

Keep performance in mind when configuring port mirroring. For example, If you mirror traffic from multiple ports, the mirrored traffic may exceed the capacity of the output interface. We recommend that you limit the amount of copied traffic by selecting specific interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter to send specific traffic to a port mirroring instance. Mirroring only the necessary packets reduces the possibility of a performance impact.

You can use port mirroring to copy any of the following:

- All packets entering or exiting an interface (in any combination)—For example, you can send copies of the packets entering some interfaces and the packets exiting other interfaces to the same local interface or VLAN. If you configure port mirroring to copy packets exiting an interface, traffic that originates on that switch or Node device (in a QFabric system) is not copied when it egresses. Only switched traffic is copied on egress. (See the limitation on egress mirroring below.)
- All packets entering a VLAN—You cannot use port mirroring to copy packets exiting a VLAN.
- Firewall-filtered sample—Sample of packets entering a port or VLAN. Configure a firewall filter to select certain packets for mirroring.



**NOTE:** Firewall filters are not supported on egress ports; therefore, you cannot specify policy-based sampling of packets exiting an interface.

## Port Mirroring Instance Types

To configure port mirroring, you configure an instance of one of the following types:

- Analyzer instance: You must specify the input and output for the instance. This instance type is useful for ensuring that all traffic transiting an interface or VLAN is mirrored and sent to the analyzer device.
- Port-mirroring instance: You do not specify an input for this instance type. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored. When you use a port-mirroring instance, you can direct traffic to it in the following ways:
  - Specify the name of the port-mirroring instance in the firewall filter using the **port-mirror-instance instance-name** action. You should use this approach if there are multiple port-mirroring instances defined.
  - Configure the filter to send the mirrored packets to the output interface defined in the instance using the **port-mirror** action. You can use this approach if there is only one port-mirroring instance defined.

## Port-Mirroring Terminology

Table 3 on page 4 lists the terms used in the documentation about port mirroring and provides definitions.

**Table 3: Port Mirroring Terms and Definitions**

Term	Description
Analyzer instance	Port-mirroring configuration that includes a name, source interfaces or source VLAN, and a destination for mirrored packets (either a local access interface or a VLAN).

Table 3: Port Mirroring Terms and Definitions (*continued*)

Port mirroring instance	A port-mirroring configuration that does not specify an input.. A firewall filter must be used to send traffic to the port mirror. Use the action <b>port-mirror-instance <i>instance-name</i></b> in the firewall filter configuration to send packets to the port mirror.
Output interface (also known as monitor interface)	<p>Access interface to which packet copies are sent and to which a device running an analyzer application is connected.</p> <p>The following limitations apply to an output interface:</p> <ul style="list-style-type: none"> <li>• Cannot also be a source port.</li> <li>• Cannot be used for switching.</li> <li>• Cannot be an aggregated Ethernet interface (LAG).</li> <li>• Does not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP).</li> <li>• Loses any existing VLAN associations when you configure it as an analyzer output interface.</li> </ul> <p>If the capacity of the output interface is insufficient to handle the traffic from the source ports, overflow packets are dropped.</p>
Output IP address	<p>IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)</p> <ul style="list-style-type: none"> <li>• An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</li> <li>• If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).</li> </ul>
Output VLAN (also known as monitor or analyzer VLAN)	<p>VLAN to which copies are sent and to which a device running an analyzer application is connected. The analyzer VLAN can span multiple switches.</p> <p>The following limitations apply to an output VLAN:</p> <ul style="list-style-type: none"> <li>• Cannot be a private VLAN or VLAN range.</li> <li>• Cannot be shared by multiple <b>analyzer</b> statements.</li> <li>• An output VLAN interface cannot be a member of any other VLAN.</li> <li>• An output VLAN interface cannot be an aggregated Ethernet interface (LAG).</li> <li>• On some switches, only one interface can be a member of the analyzer VLAN. This limitation does not apply on the QFX10000 switch if traffic is mirrored on ingress. In this case, multiple QFX10000 interfaces can belong to the output VLAN, and traffic is mirrored to all of those interfaces. If traffic is mirrored on egress on a QFX10000 switch, only one interface can be a member of the analyzer VLAN.</li> </ul>
Input interface (also known as mirrored or monitored interface)	Interface that provides traffic to be mirrored. This traffic can be entering or exiting the interface. (Ingress or egress traffic can be mirrored.) An input interface cannot also be an output interface for an analyzer.
Monitoring station	Computer running an analyzer application.
Local port mirroring	Port-mirroring configuration in which the mirrored packets are sent to an interface on the same switch.

Table 3: Port Mirroring Terms and Definitions (*continued*)

Remote port mirroring	Flooding mirrored packets to an output (analyzer) VLAN that you create to receive mirror traffic or sending the mirrored packets to a remote IP address. (You cannot send mirrored packets to a remote IP address on a QFabric system.)
Policy-based mirroring	Mirroring of packets that match the match a firewall filter term. The action <b>analyzer</b> <i>analyzer-name</i> is used in the firewall filter to send the packets to the analyzer.

## Port Mirroring and STP

The behavior of STP in a port-mirroring configuration depends on the version of Junos OS you are using:

- Junos OS 13.2X50, Junos OS 13.2X51-D25 or earlier, Junos OS 13.2X52: If you enable STP, port mirroring might not work because STP might block the mirrored packets.
- Junos OS 13.2X51-D30, Junos OS 14.1X53: STP is disabled for mirrored traffic. You must ensure that your topology prevents loops for this traffic.

## Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 6](#)
- [Remote Port Mirroring Only on page 8](#)

### Local and Remote Port Mirroring

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer** action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)
  - There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**
  - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.
- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

### Remote Port Mirroring Only

---

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnet as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.

#### Related Documentation

- [Configuring Port Mirroring on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 16](#)
- [Troubleshooting Port Mirroring on page 24](#)

## Configuring Port Mirroring

---

You use port mirroring to copy packets and send the copies to a device running an application such as a network analyzer or intrusion detection application so that you can analyze traffic without delaying it. You can mirror traffic entering or exiting a port or entering a VLAN, and you can send the copies to a local access interface or to a VLAN through a trunk interface.

We recommend that you disable port mirroring when you are not using it. To avoid creating a performance issue, if you do enable port mirroring, we recommend that you select specific input interfaces instead of using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter.



.....

**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring Port Mirroring*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

.....



.....

**NOTE:** If you want to create additional analyzers without deleting an existing analyzer, first disable the existing analyzer using the **disable analyzer analyzer-name** command.

.....





**NOTE:** You must configure port mirroring output interfaces as family **ethernet-switching**.

- [Configuring Port Mirroring for Local Analysis on page 9](#)
- [Configuring Port Mirroring for Remote Analysis on page 9](#)
- [Filtering the Traffic Entering an Analyzer on page 10](#)

## Configuring Port Mirroring for Local Analysis

To mirror interface traffic to a local interface on the switch:

1. If you want to mirror traffic that is ingressing or egressing specific interfaces, choose a name for the port-mirroring configuration and configure what traffic should be mirrored by specifying the interfaces and direction of traffic:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some VLAN packets might contain incorrect VLAN IDs.



**NOTE:** If you configure mirroring for packets that egress an access interface, the original packets lose any VLAN tags when they exit the access interface, but the mirrored (copied) packets retain the VLAN tags when they are sent to the analyzer system.

2. If you want to specify that all traffic entering a VLAN should be mirrored, choose a name for the port-mirroring configuration and specify the VLAN:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name input ingress vlan vlan-name
```



**NOTE:** You cannot configure port mirroring to copy traffic that egresses a VLAN.

3. Configure the destination interface for the mirrored packets:

```
[edit forwarding-options]
user@switch# set analyzer analyzer-name output interface interface-name
```

## Configuring Port Mirroring for Remote Analysis

To mirror traffic to a VLAN for analysis at a remote location:

1. Configure a VLAN to carry the mirrored traffic:

```
[edit]
```

```
user@switch# set vlans vlan-name vlan-id number
```

2. Configure the interface that connects to another switch (the uplink interface) to trunk mode and associate it with the appropriate VLAN:

```
[edit]
```

```
user@switch# set interfaces interface-name unit 0 family ethernet-switching port-mode trunk vlan members (vlan-name | vlan-id)
```

3. Configure the analyzer:

- a. Choose a name for the analyzer:

```
[edit forwarding-options]
```

```
user@switch# set analyzer analyzer-name
```

- b. Specify the interface to be mirrored and whether the traffic should be mirrored on ingress or egress:

```
[edit forwarding-options]
```

```
user@switch# set analyzer analyzer-name input (ingress | egress) interface interface-name
```

- c. Specify the appropriate IP address or VLAN as the output (a VLAN is specified in this example:

```
[edit forwarding-options]
```

```
user@switch# set analyzer analyzer-name output vlan (vlan-name | vlan-id)
```

If you specify an IP address as the output, note the following constraints:

- The address cannot be in the same subnet as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (**inet.0** routing table).
- The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)

## Filtering the Traffic Entering an Analyzer

In addition to specifying which traffic to mirror by configuring an analyzer, you can also use a firewall filter to exercise more control over which packets are copied. For example, you might use a filter to specify that only traffic from certain applications be mirrored. The filter can use any of the available match conditions and must have an action of modifier of **port-mirror-instance** *instance-name*. If you use the same analyzer in multiple filters or terms, the output packets are copied only once.

When you use a firewall filter as the input to a port-mirroring instance, you send the copied traffic to a local interface or a VLAN just as you do when a firewall is not involved.

To configure port mirroring with filters:

1. Configure a port-mirroring instance for local or remote analysis. Configure only the output. For example, for local analysis enter:

```
[edit forwarding-options]
```

```
user@switch# set port-mirroring-instance instance-name output interface interface-name
```



**NOTE:** You cannot configure input to this instance.

2. Create a firewall filter using any of the available match conditions. In a **then** term, specify include the action modifier **port-mirror-instance** *instance-name*.
3. Apply the firewall filter to the interfaces or VLAN that should provide the input to the analyzer:

```
[edit]
user@switch# set interfaces interface-name unit 0 family ethernet-switching filter input
filter-name
[edit]
user@switch# set vlan (vlan-name | vlan-id) filter input filter-name
```

#### Related Documentation

- [Understanding Port Mirroring on page 3](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 16](#)
- [Overview of Firewall Filters](#)

## Examples: Configuring Port Mirroring for Local Analysis

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies to a local interface for local monitoring.



**NOTE:** This example uses the Enhanced Layer 2 Software (ELS) configuration style. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

This example describes how to configure port mirroring to copy traffic sent by employee computers to a switch to an access interface on the same switch.

- [Requirements on page 11](#)
- [Overview and Topology on page 12](#)
- [Example: Mirroring All Employee Traffic for Local Analysis on page 12](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 13](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2
- A switch

## Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering interfaces on the switch to an access interface on the same switch. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

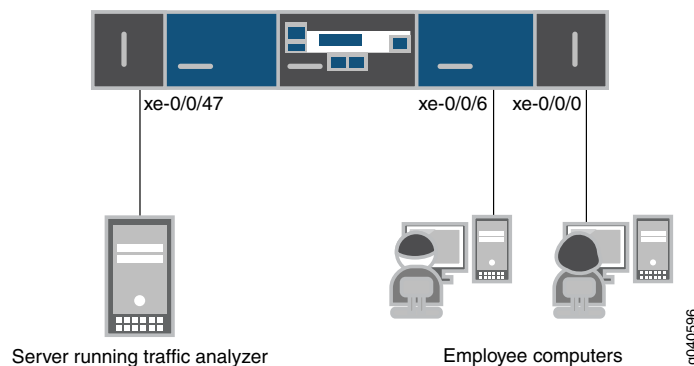
In this example, **xe-0/0/0** and **xe-0/0/6** serve as connections for employee computers. Interface **xe-0/0/47** is connected to a device running an analyzer application.



**NOTE:** Multiple ports mirrored to one interface can cause buffer overflow and dropped packets.

Figure 1 on page 12 shows the network topology for this example.

Figure 1: Network Topology for Local Port Mirroring Example



### Example: Mirroring All Employee Traffic for Local Analysis

To configure port mirroring for all traffic sent by employee computers for local analysis, perform the tasks explained in this section.

#### CLI Quick Configuration

To quickly configure local port mirroring for ingress traffic to the two ports connected to employee computers, copy the following commands and paste them into a switch terminal window:

```
[edit]
set interfaces xe-0/0/0 unit 0 family ethernet-switching
set interfaces xe-0/0/6 unit 0 family ethernet-switching
set interfaces xe-0/0/47 unit 0 family ethernet-switching
set forwarding-options analyzer employee-monitor input ingress interface xe-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface xe-0/0/6.0
set forwarding-options analyzer employee-monitor output interface xe-0/0/47.0
```

#### Step-by-Step Procedure

To configure an analyzer called **employee-monitor** and specify the input (source) interfaces and the output interface:

1. Configure the interfaces connected to employee computers as input interfaces for the port-mirror analyzer **employee-monitor**:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface xe-0/0/6.0
```

2. Configure the output analyzer interface for the **employee-monitor** analyzer. This will be the destination interface for the mirrored packets:

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor output interface xe-0/0/47.0
```

**Results** Check the results of the configuration:

```
[edit]
user@switch# show forwarding-options analyzer
employee-monitor {
  input {
    ingress {
      interface xe-0/0/0.0;
      interface xe-0/0/6.0;
    }
  }
  output {
    interface {
      xe-0/0/47.0;
    }
  }
}
```

## Example: Mirroring Employee Web Traffic with a Firewall Filter

- [Requirements on page 13](#)
- [Overview on page 13](#)
- [Configuring on page 14](#)
- [Verification on page 16](#)

### Requirements

This example uses the following hardware and software components:

- One switch
- Junos 13.2X51

### Overview

Rather than mirror all traffic, it is usually desirable to mirror only certain traffic. This is a more-efficient use of your bandwidth and hardware and might be necessary because constraints on these assets. To select specific traffic for mirroring, you use a firewall filter to match the desired traffic and direct it to a port-mirroring instance. The port-mirroring instance then copies the packets and sends them to the output VLAN, interface, or IP address.

## Configuring

To specify that the only traffic that will be mirrored is traffic sent by employees to the Web, perform the tasks explained in this section. To select this traffic for mirroring, you use a firewall filter to specify this traffic and direct it to a port-mirroring instance.

### CLI Quick Configuration

To quickly configure local port mirroring of traffic from employee computers that is destined for the Web, copy the following commands and paste them into a switch terminal window:

```
[edit]
set forwarding-options port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirror-instance employee-web-monitor
set interfaces xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee
```

### Step-by-Step Procedure

To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:

1. Configure the output interface:

```
[edit interfaces]
user@switch# set xe-0/0/47 unit 0 family ethernet-switching
```

2. Configure the **employee-web-monitor** output interface. (Configure only the output—the input comes from the filter.)

```
[edit forwarding-options]
user@switch# set port-mirroring instance employee-web-monitor output interface
xe-0/0/47.0
```

3. Configure a firewall filter called **watch-employee** that includes a term to match traffic sent to the Web and send it to the port-mirroring instance **employee-web-monitor**. Traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**) does not need to be copied, so create another term to accept that traffic before it reaches the term that sends Web traffic to the instance:

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-corp from destination-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then port-mirror-instance
employee-web-monitor
```

4. Apply the firewall filter to the appropriate interfaces as an ingress filter (egress filters do not allow analyzers):

```
[edit interfaces]
```

```

user@switch# set xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee

```

**Results** Check the results of the configuration:

```

[edit]
user@switch# show
forwarding-options {
  port-mirroring {
    instance {
      employee-web-monitor {
        output {
          interface xe-0/0/47.0;
        }
      }
    }
  }
}
...
firewall family ethernet-switching {
  filter watch-employee {
    term employee-to-web {
      from {
        destination-port 80;
      }
      then port-mirror-instance employee-web-monitor;
    }
  }
}
...
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
  xe-0/0/6 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}
}

```

## Verification

---

### *Verifying That the Analyzer Has Been Correctly Created*

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify that the port mirror analyzer has been configured as expected using the **show analyzer** command.

```
user@switch> show forwarding-options analyzer
  Port mirror name           : employee-monitor
  Mirror rate                 : 1
  Maximum packet length      : 0
  State                       : up
  Ingress monitored interfaces : xe-0/0/0.0
  Ingress monitored interfaces : xe-0/0/6.0
  Output interface           : xe-0/0/47.0
```

**Meaning** This output shows that the port-mirroring instance **employee-monitor** has a ratio of 1 (mirroring every packet, the default setting), the maximum size of the original packet that was mirrored (0 indicates the entire packet), the state of the configuration (is up indicates that the instance is mirroring the traffic entering the xe-0/0/0, and xe-0/0/6 interfaces, and sending the mirrored traffic to the xe-0/0/47 interface). If the state of the output interface is down or if the output interface is not configured, the value of state will be **down** and the instance will not be programmed for mirroring.

**Related Documentation**

- [Understanding Port Mirroring on page 3](#)

**Related Documentation**

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 8](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 16](#)

## Example: Configuring Port Mirroring for Remote Analysis

---

Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on. Port mirroring copies packets entering or exiting an interface or entering a VLAN and sends the copies either to a local interface for local monitoring or to a VLAN for remote monitoring. This example describes how to configure port mirroring for remote analysis.

- [Requirements on page 17](#)
- [Overview and Topology on page 17](#)
- [Mirroring All Employee Traffic for Remote Analysis on page 17](#)



- [Mirroring Employee-to-Web Traffic for Remote Analysis on page 18](#)
- [Verification on page 20](#)

## Requirements

This example uses the following hardware and software components:

- Junos OS Release 13.2 for the QFX Series
- A switch

## Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to an analyzer VLAN so that you can perform analysis using a remote device. The first example shows how to mirror all traffic sent by employee computers to the switch. The second example includes a filter to mirror only the employee traffic going to the Web.

In this example:

- Interfaces **ge-0/0/0** and **ge-0/0/1** are Layer 2 interfaces that connect to employee computers.
- Interface **ge-0/0/10** is a Layer 2 interface that connects to another switch.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.



**NOTE:** In addition to performing the configuration steps described here, you must also configure the analyzer VLAN (**remote-analyzer** in this example) on the other switches that are used to connect the source switch (the one in this configuration) to the one that the monitoring station is connected to.

## Mirroring All Employee Traffic for Remote Analysis

**CLI Quick Configuration** To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set forwarding-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set forwarding-options analyzer employee-monitor output vlan remote-analyzer
```

**Step-by-Step Procedure**

To configure basic remote port mirroring:

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):  

```
[edit vlans]
user@switch# set vlans remote-analyzer vlan-id 999
```
2. Configure the interface connected to another switch for trunk mode and associate it with the **remote-analyzer** VLAN:  

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```
3. Configure the **employee-monitor** analyzer:  

```
[edit forwarding-options]
user@switch# set analyzer employee-monitor
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```
4. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

**Results** Check the results of the configuration:

```
[edit]
user@switch# show
forwarding-options {
  analyzer employee-monitor {
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
    }
    output {
      vlan {
        remote-analyzer;
      }
    }
  }
}
```

## Mirroring Employee-to-Web Traffic for Remote Analysis

**CLI Quick Configuration**

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **edit** hierarchy level:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options port-mirroring instance employee-web-monitor loss-priority high output vlan 999
set firewall family ethernet-switching filter watch-employee term employee-to-web from destination-port 80
```

```

set firewall family ethernet-switching filter watch-employee term employee-to-web then
port-mirror-instance employee-web-monitor employee-web-monitor
set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```

### Step-by-Step Procedure

1. Configure the analyzer VLAN (called **remote-analyzer** in this example):  

```

[edit vlans]
user@switch# set remote-analyzer vlan-id 999

```
2. Configure an interface to associate it with the **remote-analyzer** VLAN:  

```

[edit interfaces]
user@switch# set interfaces ge-0/0/10 unit 0 family ethernet-switching interface-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999

```
3. Configure the **employee-web-monitor** analyzer. (Configure only the output—the input comes from the filter.)  

```

[edit forwarding-options]
user@switch# set forwarding-options port-mirroring instance employee-web-monitor output vlan 999

```
4. Configure a firewall filter called **watch-employee** to match traffic sent to the Web and send it to the analyzer **employee-web-monitor**:  

```

[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then port-mirror-instance employee-web-monitor

```
5. Apply the firewall filter to the appropriate interfaces as an ingress filter:  

```

[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```
6. Configure the **remote-analyzer** VLAN on the switches that connect this switch to the monitoring workstation.

**Results** Check the results of the configuration:

```

[edit]
user@switch# show
interfaces {
  ...
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        interface-mode trunk;
        vlan {
          members remote-analyzer;
        }
      }
    }
  }
}
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}

```

```
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}
...
firewall {
  family ethernet-switching {
    ...
    filter watch-employee {
      term employee-to-web {
        from {
          destination-port 80;
        }
        then port-mirror-instance employee-web-monitor;
      }
    }
  }
}
forwarding-options analyzer {
  employee-web-monitor {
    output {
      vlan {
        999;
      }
    }
  }
}
vlands {
  remote-analyzer {
    vlan-id 999;
  }
}
```

## Verification

---

### Verifying That the Analyzer Has Been Correctly Created

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
user@switch> show analyzer
Analyzer name           : employee-monitor
Output VLAN             : remote-analyzer
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
```

**Meaning** This output shows that the **employee-monitor** analyzer is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1** and is sending the mirror traffic to the analyzer **remote-analyzer**.

- Related Documentation**
- [Understanding Port Mirroring on page 3](#)
  - [Configuring Port Mirroring on page 8](#)
  - [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
  - [Overview of Firewall Filters](#)

## Example: Mirroring Employee Web Traffic with a Firewall Filter

- [Requirements on page 21](#)
- [Overview on page 21](#)
- [Configuring on page 21](#)
- [Verification on page 23](#)

### Requirements

This example uses the following hardware and software components:

- One switch
- Junos 13.2X51

### Overview

Rather than mirror all traffic, it is usually desirable to mirror only certain traffic. This is a more-efficient use of your bandwidth and hardware and might be necessary because constraints on these assets. To select specific traffic for mirroring, you use a firewall filter to match the desired traffic and direct it to a port-mirroring instance. The port-mirroring instance then copies the packets and sends them to the output VLAN, interface, or IP address.

### Configuring

To specify that the only traffic that will be mirrored is traffic sent by employees to the Web, perform the tasks explained in this section. To select this traffic for mirroring, you use a firewall filter to specify this traffic and direct it to a port-mirroring instance.

<b>CLI Quick Configuration</b>	<p>To quickly configure local port mirroring of traffic from employee computers that is destined for the Web, copy the following commands and paste them into a switch terminal window:</p> <pre>[edit] set forwarding-options port-mirroring instance employee-web-monitor output interface xe-0/0/47.0 set firewall family ethernet-switching filter watch-employee term employee-to-corp from destination-address 192.0.2.16/28 set firewall family ethernet-switching filter watch-employee term employee-to-corp from source-address 192.0.2.16/28 set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept set firewall family ethernet-switching filter watch-employee term employee-to-web from destination-port 80 set firewall family ethernet-switching filter watch-employee term employee-to-web then port-mirror-instance employee-web-monitor set interfaces xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee set interfaces xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee</pre>
<b>Step-by-Step Procedure</b>	<p>To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:</p> <ol style="list-style-type: none"> <li>1. Configure the output interface: <pre>[edit interfaces] user@switch# set xe-0/0/47 unit 0 family ethernet-switching</pre> </li> <li>2. Configure the <b>employee-web-monitor</b> output interface. (Configure only the output—the input comes from the filter.) <pre>[edit forwarding-options] user@switch# set port-mirroring instance employee-web-monitor output interface xe-0/0/47.0</pre> </li> <li>3. Configure a firewall filter called <b>watch-employee</b> that includes a term to match traffic sent to the Web and send it to the port-mirroring instance <b>employee-web-monitor</b>. Traffic to and from the corporate subnet (destination or source address of <b>192.0.2.16/28</b>) does not need to be copied, so create another term to accept that traffic before it reaches the term that sends Web traffic to the instance: <pre>[edit firewall family ethernet-switching] user@switch# set filter watch-employee term employee-to-corp from destination-address 192.0.2.16/28 user@switch# set filter watch-employee term employee-to-corp from source-address 192.0.2.16/28 user@switch# set filter watch-employee term employee-to-corp then accept user@switch# set filter watch-employee term employee-to-web from destination-port 80 user@switch# set filter watch-employee term employee-to-web then port-mirror-instance employee-web-monitor</pre> </li> <li>4. Apply the firewall filter to the appropriate interfaces as an ingress filter (egress filters do not allow analyzers): <pre>[edit interfaces] user@switch# set xe-0/0/0 unit 0 family ethernet-switching filter input watch-employee user@switch# set xe-0/0/6 unit 0 family ethernet-switching filter input watch-employee</pre> </li> </ol>
<b>Results</b>	<p>Check the results of the configuration:</p> <pre>[edit] user@switch# show</pre>

```

forwarding-options {
  port-mirroring {
    instance {
      employee-web-monitor {
        output {
          interface xe-0/0/47.0;
        }
      }
    }
  }
}
...
firewall family ethernet-switching {
  filter watch-employee {
    term employee-to-web {
      from {
        destination-port 80;
      }
      then port-mirror-instance employee-web-monitor;
    }
  }
}
...
interfaces {
  xe-0/0/0 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
  xe-0/0/6 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}

```

## Verification

### Verifying That the Analyzer Has Been Correctly Created

**Purpose** Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces and appropriate output interface.

**Action** You can verify that the port mirror analyzer has been configured as expected using the **show analyzer** command.

```
user@switch> show forwarding-options analyzer
  Port mirror name           : employee-monitor
  Mirror rate                 : 1
  Maximum packet length      : 0
  State                       : up
  Ingress monitored interfaces : xe-0/0/0.0
  Ingress monitored interfaces : xe-0/0/6.0
  Output interface           : xe-0/0/47.0
```

**Meaning** This output shows that the port-mirroring instance **employee-monitor** has a ratio of 1 (mirroring every packet, the default setting), the maximum size of the original packet that was mirrored (**0** indicates the entire packet), the state of the configuration (is up indicates that the instance is mirroring the traffic entering the xe-0/0/0, and xe-0/0/6 interfaces, and sending the mirrored traffic to the xe-0/0/47 interface). If the state of the output interface is down or if the output interface is not configured, the value of state will be **down** and the instance will not be programmed for mirroring.

**Related Documentation**

- [Understanding Port Mirroring on page 3](#)

---

## Troubleshooting Port Mirroring

- [Port Mirroring Constraints and Limitations on page 24](#)
- [Egress Port Mirroring with VLAN Translation on page 26](#)
- [Egress Port Mirroring with Private VLANs on page 26](#)

### Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 24](#)
- [Remote Port Mirroring Only on page 26](#)

---

#### Local and Remote Port Mirroring

The following constraints and limitations apply to local and remote port mirroring:

- You can create a total of four port-mirroring configurations.
- You can create a total of four port-mirroring configurations on each Node group in a QFabric system, subject to the following constraints:
  - As many as four of the configurations can be for local port mirroring.
  - As many as three of the configurations can be for remote port mirroring.
- Regardless of whether you are configuring a standalone switch or a Node group, the following limits apply:
  - There can be no more than two configurations that mirror ingress traffic. (If you configure a firewall filter to send traffic to a port mirror—that is, you use the **analyzer**



action modifier in a filter term—this counts as an ingress mirroring configuration for switch or Node group on which the filter is applied.)

- There can be no more than two configurations that mirror egress traffic.



**NOTE:** On QFabric systems, there is no system-wide limit on the total number of mirror sessions.

- You can configure no more than one type of output in one port-mirroring configuration. That is, you can use no more than one of the following to complete a **set analyzer name output** statement:
  - **interface**
  - **ip-address**
  - **vlan**
- If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs on a standalone switch or QFabric system. If you do so, some VLAN packets might contain incorrect VLAN IDs. This applies to any VLAN packets—not only the mirrored copies.
- The **ratio** and **loss-priority** options are not supported.
- Packets with physical layer errors are filtered out and are not sent to the output port or VLAN.
- If you use sFlow monitoring to sample traffic, it does not sample the mirror copies when they exit from the output interface.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis interfaces
  - Management interfaces (me0 or vme0)
  - Fibre Channel interfaces
  - Routed VLAN interfaces
- An aggregated Ethernet interface cannot be an output interface if the input is a VLAN or if traffic is sent to the analyzer by a firewall filter.
- Do not include an 802.1Q subinterface that has a unit number other than 0 in a port mirroring configuration. Port mirroring does not work with subinterfaces if their unit number is not 0. (You configure 802.1Q subinterfaces using the **vlan-tagging** statement.)
- When packet copies are sent out the output interface, they are not modified for any changes that are normally applied on egress, such as CoS rewriting.
- An interface can be the input interface for only one mirroring configuration. Do not use the same interface as the input interface for multiple mirroring configurations.
- CPU-generated packets (such as ARP, ICMP, BPDU, and LACP packets) cannot be mirrored on egress.

- VLAN-based mirroring is not supported for STP traffic.
- (QFabric systems only) If you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on different Node devices, the mirrored copies have incorrect VLAN IDs. This limitation does not apply if you configure a QFabric analyzer to mirror egress traffic and the input and output interfaces are on the *same* Node device. In this case the mirrored copies have the correct VLAN IDs (as long as you do not configure more than 2000 VLANs on the QFabric system).

### Remote Port Mirroring Only

---

The following constraints and limitations apply to remote port mirroring:

- If you configure an output IP address, the address cannot be in the same subnetwork as any of the switch's management interfaces.
- If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).
- An output VLAN cannot be a private VLAN or VLAN range.
- An output VLAN cannot be shared by multiple **analyzer** statements.
- An output VLAN interface cannot be a member of any other VLAN.
- An output VLAN interface cannot be an aggregated Ethernet interface.

### Egress Port Mirroring with VLAN Translation

**Problem Description:** If you create a port-mirroring configuration that mirrors customer VLAN (CVLAN) traffic on egress and the traffic undergoes VLAN translation before being mirrored, the VLAN translation does not apply to the mirrored packets. That is, the mirrored packets retain the service VLAN (SVLAN) tag that should be replaced by the CVLAN tag on egress. The original packets are unaffected—on these packets VLAN translation works properly, and the SVLAN tag is replaced with the CVLAN tag on egress.

**Solution** This is expected behavior.

### Egress Port Mirroring with Private VLANs

**Problem Description:** If you create a port-mirroring configuration that mirrors private VLAN (PVLAN) traffic on egress, the mirrored traffic (the traffic that is sent to the analyzer system) has the VLAN tag of the ingress VLAN instead of the egress VLAN. For example, assume the following PVLAN configuration:

- Promiscuous trunk port that carries primary VLANs pvlan100 and pvlan400.
- Isolated access port that carries secondary VLAN isolated200. This VLAN is a member of primary VLAN pvlan100.

- Community port that carries secondary VLAN comm300. This VLAN is also a member of primary VLAN pvlan100.
- Output interface (monitor interface) that connects to the analyzer system. This interface forwards the mirrored traffic to the analyzer.

If a packet for pvlan100 enters on the promiscuous trunk port and exits on the isolated access port, the original packet is untagged on egress because it is exiting on an access port. However, the mirror copy retains the tag for pvlan100 when it is sent to the analyzer.

Here is another example: If a packet for comm300 ingresses on the community port and egresses on the promiscuous trunk port, the original packet carries the tag for pvlan100 on egress, as expected. However, the mirrored copy retains the tag for comm300 when it is sent to the analyzer.

**Solution** This is expected behavior.

**Related  
Documentation**

- [Understanding Port Mirroring on page 3](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 16](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 13](#)



## PART 2

# DHCP and DHCP Relay

- [Using DHCP and DHCP Relay on page 31](#)



## CHAPTER 2

# Using DHCP and DHCP Relay

- [DHCP and BOOTP Relay Overview on page 31](#)
- [Configuring DHCP and BOOTP on page 31](#)
- [Verifying and Managing DHCP Local Server Configuration on page 32](#)
- [Verifying and Managing DHCPv6 Local Server Configuration on page 33](#)
- [Verifying and Managing DHCP Relay Configuration on page 33](#)
- [Verifying and Managing DHCPv6 Relay Configuration on page 33](#)
- [Understanding Layer 3 Logical Interfaces on page 34](#)

### DHCP and BOOTP Relay Overview

---

You can configure a Juniper Networks switch to act as a Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) relay agent. This means that if the switch receives a broadcast DHCP or BOOTP request from a locally attached host (client), it relays the message to a specified DHCP or BOOTP server. You should configure the switch to be a DHCP/BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.



**NOTE:** Because DHCP and BOOTP messages are broadcast and are not directed to a specific server, switch, or router, Juniper switches cannot function as both a DHCP server and a DHCP/BOOTP relay agent at the same time. The Junos operating system (Junos OS) generates a commit error if both options are configured at the same time, and the commit operation does not succeed until one of the options is removed.

**Related Documentation**

- [Configuring DHCP and BOOTP Relay](#)

### Configuring DHCP and BOOTP

---

You can configure a switch to act as a Dynamic Host Configuration Protocol (DHCP) and Bootstrap Protocol (BOOTP) server or DHCP relay agent. When a switch is a relay agent, if a locally attached host issues a DHCP or BOOTP request as a broadcast message, the switch relays the message to a specified DHCP or BOOTP server. You should configure

a switch to be a DHCP and BOOTP relay agent if you have locally attached hosts and a remote DHCP or BOOTP server.



**NOTE:** This task uses the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see *Configuring DHCP and BOOTP Relay*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

To configure a switch to be a server, use the `dhcp-local-server` statement. To configure a switch to be a relay agent, use the `dhcp-relay` statement.

If you want to enable BOOTP support when the switch is configured to be a DHCP server, enter the following statement:

```
[edit system services dhcp-local-server]
user@switch# set overrides bootp-support
```

If you want to enable BOOTP support when the switch is configured to be a DHCP relay agent, enter the following statement:

```
[edit forwarding-options dhcp-relay]
user@switch# set overrides bootp-support
```

## Verifying and Managing DHCP Local Server Configuration

**Purpose** View or clear information about client address bindings and statistics for the extended DHCP local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

- Action**
- To display the address bindings in the client table on the extended DHCP local server:  

```
user@host> show dhcp server binding routing-instance customer routing instance
```
  - To display extended DHCP local server statistics:  

```
user@host> show dhcp server statistics routing-instance customer routing instance
```
  - To clear the binding state of a DHCP client from the client table on the extended DHCP local server:  

```
user@host> clear dhcp server binding routing-instance customer routing instance
```
  - To clear all extended DHCP local server statistics:  

```
user@host> clear dhcp server statistics routing-instance customer routing instance
```

**Related Documentation**

- [CLI Explorer](#)



## Verifying and Managing DHCPv6 Local Server Configuration

---

**Purpose** View or clear information about client address bindings and statistics for the DHCPv6 local server.

**Action** • To display the address bindings in the client table on the DHCPv6 local server:

user@host> [show dhcpv6 server binding](#)

• To display DHCPv6 local server statistics:

user@host> [show dhcpv6 server statistics](#)

• To clear all DHCPv6 local server statistics:

user@host> [clear dhcpv6 server binding](#)

• To clear all DHCPv6 local server statistics:

user@host> [clear dhcpv6 server statistics](#)

**Related Documentation** • [CLI Explorer](#)

## Verifying and Managing DHCP Relay Configuration

---

**Purpose** View or clear address bindings or statistics for extended DHCP relay agent clients:

**Action** • To display the address bindings for extended DHCP relay agent clients:

user@host> [show dhcp relay binding](#) routing-instance *customer routing instance*

• To display extended DHCP relay agent statistics:

user@host> [show dhcp relay statistics](#) routing-instance *customer routing instance*

• To clear the binding state of DHCP relay agent clients:

user@host> [clear dhcp relay binding](#) routing-instance *customer routing instance*

• To clear all extended DHCP relay agent statistics:

user@host> [clear dhcp relay statistics](#) routing-instance *customer routing instance*

**Related Documentation** • [CLI Explorer](#)

## Verifying and Managing DHCPv6 Relay Configuration

---

**Purpose** View or clear address bindings or statistics for extended DHCPv6 relay agent clients:

**Action** • To display the address bindings for extended DHCPv6 relay agent clients:

user@host> [show dhcpv6 relay binding](#)

• To display extended DHCPv6 relay agent statistics:

user@host> [show dhcpv6 relay statistics](#)

- To clear the binding state of DHCPv6 relay agent clients:

```
user@host> clear dhcpv6 relay binding
```

- To clear all extended DHCPv6 relay agent statistics:

```
user@host> clear dhcpv6 relay statistics
```

**Related  
Documentation**

- [CLI Explorer](#)

---

## Understanding Layer 3 Logical Interfaces

---

A Layer 3 logical interface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 logical interfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks switch to a Layer 2 switch. Only one physical connection is required between the switches.

To create Layer 3 logical interfaces on a switch, enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

We recommend that you use the VLAN ID as the logical interface number when you configure the logical interface. QFX Series and EX4600 switches support a maximum of 4089 VLANs, which includes the default VLAN. You can, however, assign a VLAN ID in the range of 1 to 4094, but five of these VLAN IDs are reserved for internal use.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces.

**Related  
Documentation**

- [Interfaces Overview](#)
- [Configuring a Layer 3 Logical Interface](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

## PART 3

# Configuration Statements and Operational Commands

- [Configuration Statements \(Port Mirroring\) on page 37](#)
- [Configuration Statements \(DHCP and DHCP Relay\) on page 53](#)
- [Configuration Statements \(Encryption\) on page 73](#)
- [Operational Command \(Port Mirroring\) on page 95](#)
- [Operational Commands \(DHCP Local Server\) on page 99](#)
- [Operational Commands \(DHCP Relay Agent\) on page 135](#)



## CHAPTER 3

# Configuration Statements (Port Mirroring)

- analyzer on page 38
- egress on page 39
- ethernet-switching (Port Mirroring) on page 40
- family (Port Mirroring) on page 41
- inet (Port Mirroring) on page 42
- ingress (Port Mirroring) on page 43
- input on page 44
- instance (Port Mirroring) on page 45
- interface (Port Mirroring) on page 46
- ip-address (Port Mirroring) on page 47
- output on page 48
- port-mirroring on page 49
- routing-instance (Port Mirroring) on page 50
- vlan (Port Mirroring) on page 51

## analyzer

<b>Syntax</b>	<pre> analyzer {   name {     input {       egress {         interface (all   interface-name);         vlan (vlan-id   vlan-name);       }       ingress {         interface (all   interface-name);         vlan (vlan-id   vlan-name);       }     }     output {       interface interface-name;       ip-address ip-address;       routing-instance       vlan (vlan-id   vlan-name);     }   } } </pre>
<b>Hierarchy Level</b>	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>output vlan</b> added in Junos OS Release 12.1 for the QFX Series.</p> <p>Option <b>output ip-address</b> added in Junos OS Release 12.3 for the QFX Series for non-ELS platforms and added in 14.1X53-D10 for ELS platforms.</p>
<b>Description</b>	<p>Configure port mirroring. You can create a total of four port-mirroring configurations on the QFX Series, subject to the following limits:</p> <ul style="list-style-type: none"> <li>• There can be no more than two configurations that mirror ingress traffic.</li> <li>• There can be no more than two configurations that mirror egress traffic.</li> </ul>
<b>Default</b>	Port mirroring is disabled, and Junos OS creates no default analyzers.
<b>Options</b>	<b>all</b> —Mirror all the access interfaces. Using this option does not cause the QSFP+ or management interfaces to be mirrored.



**CAUTION:** Configuring the **all** option in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.

**name**—Name of the analyzer. The name can include as many as 125 characters; must begin with a letter; and can include uppercase letters, lowercase letters, numbers, dashes, and underscores. No other special characters are allowed.

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**

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)

## egress

**Syntax**

```
egress {
    interface (all | interface-name);
}
vlan (vlan-id | vlan-name);
```

**Hierarchy Level** For platforms without ELS:

```
[edit ethernet-switching-options analyzer name input]
```

For platforms with ELS:

```
[edit forwarding-options analyzer name input]
```

**Release Information** Statement introduced in Junos OS Release 11.2 for the QFX Series.

**Description** Specify interface or VLAN for which egressing traffic is mirrored. (The `vlan` statement is not supported on all switches.)

The statement is explained separately.



**NOTE:** If you configure Junos OS to mirror egress packets, do not configure more than 2000 VLANs. If you do so, some of the mirrored packets might contain incorrect VLAN IDs.

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

**Related Documentation**

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)

## ethernet-switching (Port Mirroring)

---

**Syntax** ethernet-switching;  
output {  
  interface *interface-name* {  
  }  
  no-filter-check;  
  }  
  vlan *vlan-name* {  
    no-tag;  
  }  
}

**Hierarchy Level** [edit forwarding-options port-mirroring [instance *name*] family]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Specify that the output interface for the port mirror will be configured as an **ethernet-switching** interface.

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

**Related Documentation**

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)



## family (Port Mirroring)

<b>Syntax</b>	<pre> family (Port Mirroring)   ethernet-switching {     output {       interface <i>interface-name</i> {       }       no-filter-check;     }     vlan <i>vlan-name</i> {       no-tag;     }   }   inet     output {       ip-address <i>address</i> {       }       routing-instance <i>instance-name</i> {         ip-address <i>address</i> {         }       }     } </pre>
<b>Hierarchy Level</b>	[edit forwarding-options port-mirroring [ <i>instance name</i> ] ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Specify the type of interface that will be used to forward port mirrored packet to an analyzer device..
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 3</a></li> <li>• <a href="#">Configuring Port Mirroring on page 8</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 11</a></li> </ul>

## inet (Port Mirroring)

**Syntax**

```
inet {
  output {
    ip-address address {
    }
    routing-instance instance-name {
      ip-address address {
      }
    }
  }
}
```

**Hierarchy Level** [edit forwarding-options port-mirroring [instance *name*] family]

**Release Information** Statement introduced in Junos OS Release 14.1X53 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Specify that the output interface will be of type **inet**. Use this statement so that you can send the mirrored packets to the IP address of the device running an analyzer application. The device can be on a remote network. When you use this feature, the mirrored packets are GRE-encapsulated. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.)



**NOTE:** An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.



**NOTE:** If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).

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

## ingress (Port Mirroring)

---

<b>Syntax</b>	<pre>ingress {   interface (all   interface-name);   vlan (vlan-id   vlan-name); }</pre>
<b>Hierarchy Level</b>	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options analyzer name input]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options analyzer name input]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Specify the interfaces or VLANs for which incoming traffic is mirrored as part of a port mirroring configuration.</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 3</a></li> <li>• <a href="#">Configuring Port Mirroring on page 8</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 11</a></li> </ul>

## input

---

<b>Syntax</b>	<pre>input {     ingress {         interface (all   <i>interface-name</i>);         vlan (<i>vlan-id</i>   <i>vlan-name</i>);     }     egress {         interface (all   <i>interface-name</i>);     } }</pre>
<b>Hierarchy Level</b>	<p>For platforms without ELS:</p> <p>[edit ethernet-switching-options <i>analyzer name</i>]</p> <p>For platforms with ELS:</p> <p>[edit forwarding-options <i>analyzer name</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	<p>Define the traffic to be mirrored. The definition can be a combination of traffic entering or exiting specific ports or VLANs.</p> <p>The statements are explained separately.</p>
<b>Default</b>	No default.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Port Mirroring on page 3</a></li><li>• <a href="#">Configuring Port Mirroring on page 8</a></li><li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 11</a></li></ul>

## instance (Port Mirroring)

```
Syntax  instance instance-name{
        family
        ethernet-switching {
            output {
                interface interface-name {
                }
                no-filter-check;
            }
            vlan vlan-name {
                no-tag;
            }
        }
        inet
        output {
            ip-address address {
            }
            routing-instance instance-name {
                ip-address address {
                }
            }
        }
    }
```

**Hierarchy Level** [edit forwarding-options port-mirroring]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.


**Description** Specify a port-mirroring configuration (instance). You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This is useful for controlling which types of traffic should be mirrored.

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



**Related Documentation**

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 13](#)

## interface (Port Mirroring)

<b>Syntax</b>	interface (all   <i>interface-name</i> );
<b>Hierarchy Level</b>	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <i>analyzer name</i> input (egress   ingress)],</pre> <pre>[edit ethernet-switching-options <i>analyzer name</i> output]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options <i>analyzer name</i> input (egress   ingress)]</pre> <pre>[edit forwarding-options <i>analyzer name</i> output]</pre> <pre>[edit forwarding-options port-mirroring[instance <i>name</i>] family ethernet-switching <i>output</i>]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Specify the interfaces for which ingressing traffic is mirrored. Specify the interface that mirrored traffic should be copied to (the output interface).
<b>Options</b>	<p>all—Apply port mirroring to all interfaces on the switch (except the output interface). Mirroring a high volume of traffic can cause performance issues, so you should generally select specific input interfaces.</p>
<div>  <p><b>CAUTION:</b> Configuring <b>all</b> in a QFabric system causes all the access interfaces on all the nodes to be mirrored. Be cautious about using this option on a QFabric system.</p> </div>	
<p><i>interface-name</i>—Apply port mirroring to the specified interface only.</p>	
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 3</a></li> <li>• <a href="#">Configuring Port Mirroring on page 8</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 11</a></li> </ul>

## ip-address (Port Mirroring)

<b>Syntax</b>	<code>ip-address <i>ip-address</i>;</code>
<b>Hierarchy Level</b>	<code>[edit forwarding-options] <b>analyzer name</b> <b>output</b></code> <code>[edit forwarding-options port-mirroring [instance <i>name</i>] family ethernet-switching <b>output</b></code> <code><b>interface name</b>]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the IP address to which traffic should be mirrored (the IP address of the analyzer system). The device can be on a remote network. The analyzer device must be able to de-encapsulate GRE-encapsulated packets, or the GRE-encapsulated packets must be de-encapsulated before reaching the analyzer device. (You can use a network sniffer to de-encapsulate the packets.) This statement is not supported on QFabric systems.
<div>  <p><b>NOTE:</b> An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</p> </div>	
<div>  <p><b>NOTE:</b> If you create virtual routing instances and also create an analyzer configuration that includes an output IP address, the output address belongs to the default virtual routing instance (inet.0 routing table).</p> </div>	
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## output

---

<b>Syntax</b>	<pre>output {     interface <i>interface-name</i>;     ip-address <i>ip-address</i>;     vlan (<i>vlan-id</i>   <i>vlan-name</i>);     routing-instance <i>instance-name</i> {         ip-address <i>address</i> {         }     } }</pre>
<b>Hierarchy Level</b>	<p>For platforms without ELS:</p> <p>[edit ethernet-switching-options <i>analyzer name</i>]</p> <p>For platforms with ELS:</p> <p>[edit forwarding-options <i>analyzer name</i>]</p> <p>[edit forwarding-options port-mirroring [instance <i>name</i>] family ethernet-switching ]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>output vlan</b> added in Junos OS Release 12.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure the destination for mirrored traffic, either an interface on the switch (for local monitoring) or a VLAN (for remote monitoring).</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>



## port-mirroring

```

Syntax  port-mirroring {
        family {
            ethernet-switching
            output {
                interface interface-name {
                }
                no-filter-check;
            }
            vlan vlan-name {
                no-tag;
            }
        }
        inet
        output {
            ip-address address {
            }
            routing-instance instance-name {
                ip-address address {
                }
            }
        }
    }
    instance instance-name {
        family (Port Mirroring)
        ethernet-switching {
            output {
                interface interface-name {
                }
                no-filter-check;
            }
            vlan vlan-name {
                no-tag;
            }
        }
        inet
        output {
            ip-address address {
            }
            routing-instance instance-name {
                ip-address address {
                }
            }
        }
    }
}

```

**Hierarchy Level** [edit forwarding-options ]

**Release Information** Statement introduced in Junos OS Release 13.2 for the QFX Series.

**Description** Create a port-mirroring configuration.

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

**Related Documentation**

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 11](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 13](#)

---

## routing-instance (Port Mirroring)

---

**Syntax** routing-instance *instance-name*;

**Hierarchy Level** [edit forwarding-options] [analyzer name](#) [output](#)]  
[edit forwarding-options port-mirroring [instance *name*] family inet [output interface name](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure a port mirroring instance. You do not specify an input for this instance. Instead, you, create a firewall filter that specifies the required traffic and directs it to the mirror. This instance type is useful for controlling which types of traffic should be mirrored.

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

## vlan (Port Mirroring)

<b>Syntax</b>	<code>vlan (vlan-id   vlan-name) { no-tag;</code>
<b>Hierarchy Level</b>	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options analyzer name input ingress], [edit ethernet-switching-options analyzer name output]</pre> <p>For platforms with ELS:</p> <pre>[edit forwarding-options analyzer name input (egress   ingress)] [edit forwarding-options analyzer name output] [edit forwarding-options port-mirroring[instance name] family ethernet-switching output]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option <b>output</b> <b>vlan</b> added in Junos OS Release 12.1 for the QFX Series.</p> <p>Option <b>no-tag</b> added in Junos OS Release 15.1X53-D10 for QFX10000 switches.</p>
<b>Description</b>	<p>When used in an <b>input</b> statement, specifies that traffic entering or exiting a VLAN should be mirrored. (You can include this statement in an <b>ingress</b> statement or <b>egress</b> statement within the <b>input</b> statement. It is not supported in an <b>egress</b> statement on all switches)</p> <p>When used in an <b>output</b> statement, specifies that mirrored traffic to be sent to a VLAN for remote monitoring.</p> <p>On some switches, only one interface can be a member of the output (analyzer) VLAN. This limitation does not apply on the QFX10000 switch if traffic is mirrored on ingress. In this case, multiple QFX10000 interfaces can belong to the output VLAN, and traffic is mirrored to all of those interfaces. If traffic is mirrored on egress on a QFX10000 switch, only one interface can be a member of the analyzer VLAN.</p>
<b>Options</b>	<p><b>vlan-id</b>—Numeric VLAN identifier.</p> <p><b>vlan-name</b>—Name of the VLAN.</p> <p><b>no-tag</b>—Specifies that remote mirrored packets are not tagged with the tag of the output (analyzer) VLAN.</p>
<b>Required Privilege Level</b>	<p><b>routing</b>—To view this statement in the configuration.</p> <p><b>routing-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Port Mirroring on page 3</a></li> <li>• <a href="#">Configuring Port Mirroring on page 8</a></li> <li>• <a href="#">Examples: Configuring Port Mirroring for Local Analysis on page 11</a></li> </ul>



## CHAPTER 4

# Configuration Statements (DHCP and DHCP Relay)

- [dhcp-local-server](#) on page 54
- [dhcp-relay](#) on page 60
- [dhcp-service](#) on page 65
- [forwarding-options](#) on page 66
- [forward-only \(DHCP Relay Agent\)](#) on page 71

## dhcp-local-server

```

Syntax  dhcp-local-server {
        access-profile profile-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
    }
    dhcpv6 {
        access-profile profile-name;
        authentication {
            ...
        }
        group group-name {
            access-profile profile-name;
            authentication {
                ...
            }
        }
        interface interface-name {
            access-profile profile-name;
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
                method {
                    bfd {
                        version (0 | 1 | automatic);
                        minimum-interval milliseconds;
                        minimum-receive-interval milliseconds;
                        multiplier number;
                        no-adaptation;
                        transmit-interval {
                            minimum-interval milliseconds;
                            threshold milliseconds;
                        }
                    }
                    detection-time {
                        threshold milliseconds;
                    }
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
        }
    }
    overrides {

```

```

        interface-client-limit number;
        multi-address-embedded-option-response;
        process-inform {
            pool pool-name;
        }
        rapid-commit;
    }
    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    delegated-pool;
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    delegated-pool;
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    multi-address-embedded-option-response;
    process-inform {
        pool pool-name;
    }
    rapid-commit;
}

```

```

}
reconfigure {
  attempts attempt-count;
  clear-on-abort;
  strict;
  timeout timeout-value;
  token token-value;
  trigger {
    radius-disconnect;
  }
}
route-suppression;
service-profile dynamic-profile-name;
}
duplicate-clients-in-subnet (incoming-interface | option-82);
dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
  primary-profile-name>;
forward-snooped-clients (all-interfaces | configured-interfaces |
  non-configured-interfaces);
group group-name {
  authentication {
    ...
  }
  dynamic-profile profile-name <aggregate-clients (merge | replace) | use-primary
    primary-profile-name>;
  interface interface-name {
    exclude;
    liveness-detection {
      failure-action (clear-binding | clear-binding-if-interface-up | log-only);
      method {
        bfd {
          version (0 | 1 | automatic);
          minimum-interval milliseconds;
          minimum-receive-interval milliseconds;
          multiplier number;
          no-adaptation;
          transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
          }
          detection-time {
            threshold milliseconds;
          }
          session-mode (automatic | multihop | singlehop);
          holddown-interval milliseconds;
        }
      }
    }
  }
}
overrides {
  client-discover-match (option60-and-option82 | incoming-interface);
  include-option-82 {
    forcerenew;
    nak;
  }
  interface-client-limit number;
  process-inform {

```



```

        pool pool-name;
    }
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
overrides {
    client-discover-match (option60-and-option82 | incoming-interface);
    include-option-82 {
        forcerenew;
        nak;
    }
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}
requested-ip-network-match subnet-mask
route-suppression;
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    client-discover-match <option60-and-option82 | incoming-interface>;
    interface-client-limit number;
    process-inform {
        pool pool-name;
    }
}
pool-match-order {
    external-authority;
    ip-address-first;
    option-82;
}
reconfigure {

```

```

    attempts attempt-count;
    clear-on-abort;
    strict;
    timeout timeout-value;
    token token-value;
    trigger {
        radius-disconnect;
    }
}
requested-ip-network-match subnet-mask;
route-suppression;
service-profile dynamic-profile-name;
}

```

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

**Release Information** Statement introduced in Junos OS Release 9.0.  
 Statement introduced in Junos OS Release 12.1 for EX Series switches.  
 Statement introduced in Junos OS Release 13.2X51 for the QFX Series.  
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure Dynamic Host Configuration Protocol (DHCP) local server options on the router or switch and enable the router or switch to function as an extended DHCP local server. The DHCP local server receives DHCP request and reply packets from DHCP clients and then responds with an IP address and other optional configuration information to the client.

The extended DHCP local server is incompatible with the DHCP server on J Series routers and so is not supported on J Series routers. Also, the DHCP local server and the DHCP/BOOTP relay server, which are configured under the **[edit forwarding-options helpers]** hierarchy level, cannot both be enabled on the router or switch at the same time. The extended DHCP local server is fully compatible with the extended DHCP relay feature.

The **dhcpv6** stanza configures the router or switch to support Dynamic Host Configuration Protocol for IPv6 (DHCPv6). The DHCPv6 local server is fully compatible with the extended DHCP local server and the extended DHCP relay feature.



**NOTE:** When you configure the **dhcp-local-server** statement at the routing instance hierarchy level, you must use a routing instance type of **virtual-router**.

The remaining statements are explained separately.

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

- Related Documentation**
- *Extended DHCP Local Server Overview*
  - *DHCPv6 Local Server Overview*

## dhcp-relay

```

Syntax  dhcp-relay {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
        dhcpv6 {
            active-server-group server-group-name;
            authentication {
                password password-string;
                username-include {
                    circuit-type;
                    client-id;
                    delimiter delimiter-character;
                    domain-name domain-name-string;
                    interface-name;
                    logical-system-name;
                    relay-agent-interface-id;
                    relay-agent-remote-id;
                    relay-agent-subscriber-id;
                    routing-instance-name;
                    user-prefix user-prefix-string;
                }
            }
            dynamic-profile profile-name {
                aggregate-clients (merge | replace);
                use-primary primary-profile-name;
            }
        }
    }
    group group-name {
        active-server-group server-group-name;
        authentication {
            ...
        }
        dynamic-profile profile-name {
            ...
        }
        interface interface-name {
            exclude;
            liveness-detection {
                failure-action (clear-binding | clear-binding-if-interface-up | log-only);
            }
        }
    }

```

```

method {
  bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    detection-time {
      threshold milliseconds;
    }
    session-mode (automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
}
}
overrides {
  ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
  ...
}
relay-agent-interface-id {
  ...
}
service-profile dynamic-profile-name;
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode (automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
}
}

```

```

overrides {
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
    non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
}
dynamic-profile profile-name {
    ...
}
interface interface-name {
    exclude;
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
                session-mode (automatic | multihop | singlehop);
                holddown-interval milliseconds;
            }
        }
    }
}
overrides {
    ...
}

```

```

    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    ...
}
relay-option-82 {
    ...
}
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}
server-group {
    server-group-name {

```

```
        server-ip-address;  
    }  
}  
service-profile dynamic-profile-name;  
}
```

**Hierarchy Level** [edit forwarding-options],  
[edit vlans forwarding-options]

**Release Information** Statement introduced in Junos OS Release 11.3 for the QFX Series.

**Description** Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the switch and enable the switch to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

DHCP relay supports the attachment of dynamic profiles and also interacts with the local AAA Service Framework to use back-end authentication servers, such as RADIUS, to provide subscriber authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.

The extended DHCP and DHCPv6 relay agent options configured with the **dhcp-relay** and **dhcpv6** statements are incompatible with the DHCP/BOOTP relay agent options configured with the **bootp** statement. As a result, the extended DHCP or DHCPv6 relay agent and the DHCP/BOOTP relay agent cannot both be enabled on the router at the same time.

The remaining statements are explained separately.

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

**Related Documentation**

- [Configuring DHCP and BOOTP on page 31](#)



---

## dhcp-service

---

<b>Syntax</b>	<pre>dhcp-service {     dhcp-snooping-file (<i>local_pathname</i>   <i>remote_URL</i>);     write-interval <i>interval</i>; }</pre>
<b>Hierarchy Level</b>	[edit system processes]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1 for MX Series routers.
<b>Description</b>	<p>Enable DHCP services on the device. DHCP services automate network-parameter assignment to network devices. The DHCP service process is enabled by default. However, by default, IP-MAC bindings in the DHCP snooping database do not persist through device reboots. You can improve performance after rebooting by configuring the IP-MAC bindings to persist, by configuring a storage location for the DHCP database file. When specifying the location for the DHCP database, you must also specify how frequently the switch writes the database entries into the DHCP snooping database file.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Persistent Bindings in the DHCP or DHCPv6 Snooping Database to Improve Network Performance (CLI Procedure)</i></li></ul>

## forwarding-options

```
Syntax forwarding-options {
    dhcp-relay {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                mac-address;
                option-60;
                option-82 <circuit-id> <remote-id>;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
    }
    dhcpv6 {
        active-server-group server-group-name;
        authentication {
            password password-string;
            username-include {
                circuit-type;
                client-id;
                delimiter delimiter-character;
                domain-name domain-name-string;
                interface-name;
                logical-system-name;
                relay-agent-interface-id;
                relay-agent-remote-id;
                relay-agent-subscriber-id;
                routing-instance-name;
                user-prefix user-prefix-string;
            }
        }
    }
    dynamic-profile profile-name {
        aggregate-clients (merge | replace);
        use-primary primary-profile-name;
    }
    group group-name {
        active-server-group server-group-name;
        authentication {
            ...
        }
        dynamic-profile profile-name {
            ...
        }
    }
    interface interface-name {
        exclude;
        liveness-detection {
            failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        }
    }
}
```

```

method {
  bfd {
    version (0 | 1 | automatic);
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    detection-time {
      threshold milliseconds;
    }
    session-mode(automatic | multihop | singlehop);
    holddown-interval milliseconds;
  }
}
overrides {
  ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
service-profile dynamic-profile-name;
}
overrides {
  ...
}
relay-agent-interface-id {
  ...
}
service-profile dynamic-profile-name;
liveness-detection {
  failure-action (clear-binding | clear-binding-if-interface-up | log-only);
  method {
    bfd {
      version (0 | 1 | automatic);
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      detection-time {
        threshold milliseconds;
      }
      session-mode(automatic | multihop | singlehop);
      holddown-interval milliseconds;
    }
  }
}
}

```

```

overrides {
    allow-snooped-clients;
    interface-client-limit number;
    no-allow-snooped-clients;
    no-bind-on-request;
    send-release-on-delete;
}
relay-agent-interface-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
dynamic-profile profile-name {
    aggregate-clients (merge | replace);
    use-primary primary-profile-name;
}
forward-snooped-clients (all-interfaces | configured-interfaces |
non-configured-interfaces);
group group-name {
    active-server-group server-group-name;
    authentication {
        ...
    }
}
dynamic-profile profile-name {
    ...
}
interface interface-name {
    exclude;
    liveness-detection {
        failure-action (clear-binding | clear-binding-if-interface-up | log-only);
        method {
            bfd {
                version (0 | 1 | automatic);
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                multiplier number;
                no-adaptation;
                transmit-interval {
                    minimum-interval milliseconds;
                    threshold milliseconds;
                }
                detection-time {
                    threshold milliseconds;
                }
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    ...
}

```

```

    service-profile dynamic-profile-name;
    trace;
    upto upto-interface-name;
}
overrides {
    ...
}
relay-option-60 {
    ...
}
relay-option-82 {
    ...
}
service-profile dynamic-profile-name;
}
liveness-detection {
    failure-action (clear-binding | clear-binding-if-interface-up | log-only);
    method {
        bfd {
            version (0 | 1 | automatic);
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            detection-time {
                threshold milliseconds;
            }
            session-mode (automatic | multihop | singlehop);
            holddown-interval milliseconds;
        }
    }
}
overrides {
    allow-snooped-clients;
    always-write-giaddr;
    always-write-option-82;
    client-discover-match <option60-and-option82>;
    disable-relay;
    interface-client-limit number;
    layer2-unicast-replies;
    no-allow-snooped-clients;
    no-bind-on-request;
    proxy-mode;
    replace-ip-source-with;
    send-release-on-delete;
    trust-option-82;
}
relay-option-82 {
    circuit-id {
        prefix prefix;
        use-interface-description (logical | device);
    }
}

```

```
    }
    server-group {
        server-group-name {
            server-ip-address;
        }
    }
    service-profile dynamic-profile-name;
}
dhcp-security {
    arp-inspection;
    group group-name {
        interface interface-name {
            static-ip ip-address {
                mac mac-address;
            }
        }
    }
    overrides {
        no-option82;
        trusted;
        untrusted;
    }
}
ip-source-guard;
no-dhcp-snooping;
option-82 {
    circuit-id {
        prefix {
            host-name;
            logical-system-name;
            routing-instance-name;
        }
        use-interface-description (device | logical);
        use-vlan-id;
    }
    remote-id {
        host-name hostname;
        use-interface-description (device | logical);
        use-string string;
    }
    vendor-id {
        use-string string;
    }
}
}
fip-security {
    examine-vn2vf;
    examine-vn2vn {
        beacon-period milliseconds;
    }
    fc-map fc-map-value;
    interface interface-name {
        (fcoe-trusted | no-fcoe-trusted;)
    }
}
}
```

<b>Hierarchy Level</b>	[edit] [edit vlans]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.3 for QFX Series switches.
<b>Description</b>	Configure traffic forwarding.  The statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

### forward-only (DHCP Relay Agent)

<b>Syntax</b>	forward-only { routing-instance <current   default   <i>routing-instance-name</i> >; }
<b>Hierarchy Level</b>	[edit forwarding-options dhcp-relay], [edit forwarding-options dhcp-relay dhcpv6], [edit forwarding-options dhcp-relay group <i>group-name</i> ], [edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> forwarding-options dhcp-relay ...]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D25 for QFX Series.
<b>Description</b>	Specify the location of the DHCP server when configuring secure DHCP traffic between the DHCP server and DHCP client when the two reside in different virtual routing instances.
<b>Default</b>	Routing instance from where the configuration is applied.
<b>Options</b>	<b>routing-instance</b> —(Optional) Routing instance in which the DHCP server resides. <ul style="list-style-type: none"> <li><b>current</b>—Routing instance from which the configuration is applied.</li> <li><b>default</b>—Root routing instance.</li> </ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs</i></li> <li><i>Configuring DHCP Message Exchange Between DHCP Server and Clients in Different Virtual Routing Instances</i></li> </ul>





## CHAPTER 5

# Configuration Statements (Encryption)

- [authentication-key-chains on page 74](#)
- [cache-size on page 75](#)
- [cache-timeout-negative on page 76](#)
- [ca-name on page 77](#)
- [certificates on page 78](#)
- [certification-authority on page 79](#)
- [crl \(Encryption Interface\) on page 80](#)
- [encoding on page 80](#)
- [enrollment-retry on page 81](#)
- [enrollment-url on page 81](#)
- [file on page 82](#)
- [key \(Authentication Keychain\) on page 83](#)
- [key-chain \(Security\) on page 84](#)
- [ldap-url on page 85](#)
- [local on page 86](#)
- [maximum-certificates on page 87](#)
- [path-length on page 87](#)
- [secret on page 88](#)
- [security on page 89](#)
- [ssh-known-hosts on page 90](#)
- [start-time \(Authentication Key Transmission\) on page 91](#)
- [traceoptions on page 92](#)

## authentication-key-chains

<b>Syntax</b>	<pre> authentication-key-chains {   key-chain key-chain-name {     description text-string;     key key {       algorithm (md5   hmac-sha-1);       options (basic   isis-enhanced);       secret secret-data;       start-time yyyy-mm-dd.hh:mm:ss;     }     tolerance seconds;   } } </pre>
<b>Hierarchy Level</b>	[edit security]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in JUNOS OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure authentication key updates for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol. When the <b>authentication-key-chains</b> statement is configured at the <b>[edit security]</b> hierarchy level, and is associated with the BGP, LDP, or IS-IS protocols at the <b>[edit protocols]</b> hierarchy level or with the BFD protocol using the <b>bfd-liveness-detection</b> statement, authentication key updates can occur without interrupting routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li> <li>• <i>Example: Configuring BFD Authentication for Securing Static Routes</i></li> <li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li> </ul>

## cache-size

---

<b>Syntax</b>	cache-size <i>bytes</i> ;
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the cache size for digital certificates.
<b>Options</b>	<b>bytes</b> —Cache size for digital certificates. <b>Range:</b> 64 through 4,294,967,295 <b>Default:</b> 2 megabytes (MB)



**NOTE:** We recommend that you limit your cache size to 4 MB.

---

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## cache-timeout-negative

---

<b>Syntax</b>	cache-timeout-negative <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure a negative cache for digital certificates.
<b>Options</b>	<b>seconds</b> —Negative time to cache digital certificates, in seconds. <b>Range:</b> 10 through 4,294,967,295 <b>Default:</b> 20



**CAUTION:** Configuring a large negative cache value can lead to a denial-of-service attack.

---

<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

---

## ca-name

---

<b>Syntax</b>	<code>ca-name <i>ca-identity</i>;</code>
<b>Hierarchy Level</b>	[edit security certificates <a href="#">certification-authority</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the certificate authority (CA) identity to use in the certificate request.
<b>Options</b>	<i>ca-identity</i> —CA identity to use in the certificate request.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

## certificates

<b>Syntax</b>	<pre> certificates {   cache-size bytes;   cache-timeout-negative seconds;   certification-authority ca-profile-name {     ca-name ca-identity;     crt file-name;     encoding (binary   pem);     enrollment-url url-name;     file certificate-filename;     ldap-url url-name;   }   enrollment-retry attempts;   local certificate-name {     certificate-key-string;     load-key-file URL filename;   }   maximum-certificates number;   path-length certificate-path-length; } </pre>
<b>Hierarchy Level</b>	[edit security]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>(Encryption interface on M Series and T Series routers and EX Series switches only)</p> <p>Configure the digital certificates for IPsec.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## certification-authority

<b>Syntax</b>	<pre>certification-authority <i>ca-profile-name</i> {   <i>ca-name</i> <i>ca-identity</i>;   <i>crl</i> <i>file-name</i>;   <i>encoding</i> (binary   pem);   <i>enrollment-url</i> <i>url-name</i>;   <i>file</i> <i>certificate-filename</i>;   <i>ldap-url</i> <i>url-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced before Junos OS Release 12.1 for the SRX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure certification authority (CA) for X.509 certificate.</p> <p>(Encryption interface on M Series and T Series routers and EX Series switches only)</p> <p>Configure a certificate authority profile name.</p>
<b>Options</b>	<ul style="list-style-type: none"> <li>• <i>profile-name</i>—Name of this CA configuration.</li> <li>• <i>ca-name</i> <i>ca-identity</i>—Name of the CA.</li> <li>• <i>crl</i> <i>file-name</i>—Certificate revocation list (CRL) filename.</li> <li>• <i>encoding</i>—Certificate encoding, either <b>binary</b> or <b>pem</b> (privacy-enhanced mail).</li> <li>• <i>enrollment-url</i> <i>url-name</i>—Enrollment URL.</li> <li>• <i>file</i> <i>certificate-filename</i>—Certificate filename.</li> <li>• <i>ldap-url</i> <i>url-name</i>—Lightweight Directory Access Protocol (LDAP) URL.</li> </ul>
<b>Required Privilege Level</b>	<p><b>admin</b>—To view this statement in the configuration.</p> <p><b>admin-control</b>—To add this statement to the configuration</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## crl (Encryption Interface)

---

<b>Syntax</b>	<code>crl file-name;</code>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.
<b>Options</b>	<b>file-name</b> —Specify the file from which to read the CRL.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

## encoding

---

<b>Syntax</b>	<code>encoding (binary   pem);</code>
<b>Hierarchy Level</b>	[edit security ike policy <i>ike-peer-address</i> ], [edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the file format used for the <b>local-certificate</b> and <b>local-key-pair</b> statements.
<b>Options</b>	<b>binary</b> —Binary file format.  <b>pem</b> —Privacy-enhanced mail (PEM), an ASCII base 64 encoded format. <b>Default:</b> <b>binary</b>
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li><li>• <i>Configuring an IKE Policy for Digital Certificates for an ES PIC</i></li></ul>



## enrollment-retry

<b>Syntax</b>	<code>enrollment-retry <i>attempts</i>;</code>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify how many times a router or switch can resend a digital certificate request.
<b>Options</b>	<b><i>attempts</i></b> —Number of enrollment retries. <b>Range:</b> 0 through 100 <b>Default:</b> 0
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## enrollment-url

<b>Syntax</b>	<code>enrollment-url <i>url-name</i>;</code>
<b>Hierarchy Level</b>	[edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify where your router or switch sends Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests (certificate authority URL).
<b>Options</b>	<b><i>url-name</i></b> —Certificate authority URL.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## file

---

<b>Syntax</b>	<code>file <i>certificate-filename</i>;</code>
<b>Hierarchy Level</b>	[edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Specify the file from which to read the digital certificate.
<b>Options</b>	<i>certificate-filename</i> —File from which to read the digital certificate.
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Digital Certificates for an ES PIC</i></li></ul>

## key (Authentication Keychain)

<b>Syntax</b>	<pre>key key {   algorithm (md5   hmac-sha-1);   options (basic   isis-enhanced);   secret secret-data;   start-time yyyy-mm-dd.hh:mm:ss; }</pre>
<b>Hierarchy Level</b>	[edit security authentication-key-chains key-chain <i>key-chain-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Configure the authentication element.
<b>Options</b>	<p><b>key</b>—Each key within a keychain is identified by a unique integer value.</p> <p><b>Range:</b> 0 through 63</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li> <li>• <i>Example: Configuring BFD Authentication for Securing Static Routes</i></li> <li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li> </ul>

## key-chain (Security)

---


<b>Syntax</b>	<pre>keychain <i>key-chain-name</i> {     description <i>text-string</i>;     key <i>key</i> {         algorithm (md5   hmac-sha-1);         options (basic   isis-enhanced);         secret <i>secret-data</i>;         start-time <i>yyyy-mm-dd.hh:mm:ss</i>;     }     tolerance <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit security authentication-key-chains]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Create the key-chain configuration for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.
<b>Options</b>	<b><i>key-chain-name</i></b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">authentication-key-chains on page 74</a></li><li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li><li>• <i>Example: Configuring BFD Authentication for Securing Static Routes</i></li><li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li></ul>

## ldap-url

---

<b>Syntax</b>	<ldap-url <i>url-name</i> >;
<b>Hierarchy Level</b>	[edit security certificates <a href="#">certification-authority</a> <i>ca-profile-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series,
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) (Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.
<b>Options</b>	<i>url-name</i> —Name of the LDAP URL.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## local

<b>Syntax</b>	<pre>local <i>certificate-name</i> {     <i>certificate-key-string</i>;     load-key-file <i>URL filename</i>; }</pre>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Import a paired X.509 private key and authentication certificate, to enable Junos XML protocol client applications to establish Secure Sockets Layer (SSL) connections to the router or switch.
<div>  <p><b>NOTE:</b> For FIPS mode, the digital security certificates must be compliant with the National Institute of Standards and Technology (NIST) SP 800-131A standard.</p> </div>	
<b>Options</b>	<p><b><i>certificate-name</i></b>—Name that uniquely identifies the certificate.</p> <p><b><i>certificate-key-string</i></b>—String of alphanumeric characters that constitute the private key and certificate.</p> <p><b><i>load-key-file URL filename</i></b>—File that contains the private key and certificate. It can be one of two types of values:</p> <ul style="list-style-type: none"> <li>• Pathname of a file on the local disk (assuming you have already used another method to copy the certificate file to the router's or switch's local disk)</li> <li>• URL to the certificate file location (for instance, on the computer where the Junos XML protocol client application runs)</li> </ul>
<b>Required Privilege Level</b>	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Importing SSL Certificates for Junos XML Protocol Support</i></li> </ul>

## maximum-certificates

<b>Syntax</b>	<code>maximum-certificates <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the maximum number of peer digital certificates to be cached.
<b>Options</b>	<b><i>number</i></b> —Maximum number of peer digital certificates to be cached. <b>Range:</b> 64 through 4,294,967,295 peer certificates <b>Default:</b> 1024 peer certificates
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## path-length

<b>Syntax</b>	<code>path-length <i>certificate-path-length</i>;</code>
<b>Hierarchy Level</b>	[edit security <a href="#">certificates</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Configure the digital certificate path length.
<b>Options</b>	<b><i>certificate-path-length</i></b> —Digital certificate path length. <b>Range:</b> 2 through 15 certificates <b>Default:</b> 15 certificates
<b>Required Privilege Level</b>	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Digital Certificates for an ES PIC</i></li> </ul>

## secret

---

<b>Syntax</b>	<code>secret <i>secret-data</i>;</code>
<b>Hierarchy Level</b>	[edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Specify a password in encrypted text or plain text format. The secret password always appears in encrypted format.
<b>Options</b>	<b><i>secret-data</i></b> —Password to use; it can include spaces if the character string is enclosed in quotation marks.
<b>Required Privilege Level</b>	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li><li>• <i>Example: Configuring BFD Authentication for Securing Static Routes</i></li><li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li></ul>



## security

```
Syntax  security {
    authentication-key-chains {
        key-chain key-chain-name {
            key key {
                secret secret-data;
                start-time yyyy-mm-dd.hh:mm:ss;
            }
        }
    }
    certificates {
        cache-size bytes;
        cache-timeout-negative seconds;
        certification-authority ca-profile-name {
            ca-name ca-identity;
            crl file-name;
            encoding (binary | pem);
            enrollment-url url-name;
            file certificate-filename;
            ldap-url url-name;
        }
        enrollment-retry attempts;
        local certificate-filename {
            certificate-key-string;
            load-key-file key-file-name;
        }
        maximum-certificates number;
        path-length certificate-path-length;
    }
    ssh-known-hosts {
        host {
            fetch-from-server host-name;
            load-key-file file-name;
        }
    }
    traceoptions {
        file filename <files number> <size size>;
        flag flag;
        level level;
        no-remote-trace
    }
}
```

Hierarchy Level [edit]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure security services. Most of the configuration statements do not have default values. If you do not specify an identifier for a statement that does not have a default value, you cannot commit the configuration.

**Required Privilege  
Level**

**Related  
Documentation**

## ssh-known-hosts

---

**Syntax**    ssh-known-hosts {  
              host *host-name* {  
                  fetch-from-server *host-name*;  
                  load-key-file *file-name*;  
              }  
              }

**Hierarchy Level**    [edit security ssh-known-hosts]

**Release Information**    Statement introduced in Junos OS Release 11.1 for the QFX Series.  
                              Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description**    Configure SSH support for known hosts and for administering SSH host key updates.

**Options**    **host *host-name***—Hostname of the SSH known host entry. This option has the following suboptions:

- **fetch-from-server *host-name***—Retrieve SSH public host key information from a specified server.
- **load-key-file *filename***—Import SSH host key information from the `/var/tmp/ssh-known-hosts` file.


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

**Related  
Documentation**    • *Understanding Security Features on the QFabric System*  
                      • *Configuring SSH Host Keys for Secure Copying of Data*

## start-time (Authentication Key Transmission)

<b>Syntax</b>	<code>start-time (now   yyyy-mm-dd.hh:mm:ss);</code>
<b>Hierarchy Level</b>	[edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6.</p> <p>Support for the BFD protocol introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Support for IS-IS introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 11.3 for QFX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Specify a start time for key transmission. You do not need to specify an end time for the key. If a new key is present with a new start time, the keychain rolls over to the new one. The start time must be unique within the keychain.</p>
<b>Options</b>	<p><b>now</b>—Start time as the current year, month, day, hour, minute, and second.</p> <p><b>yyyy-mm-dd.hh:mm:ss</b>—Start time in UTC (Coordinated Universal Time). The start time must be unique within the keychain.</p>
<b>Required Privilege Level</b>	<p><b>admin</b>—To view this statement in the configuration.</p> <p><b>admin-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i></li> <li>• <i>Example: Configuring BFD Authentication for Securing Static Routes</i></li> <li>• <i>Example: Configuring BFD Authentication for Securing Static Routes</i></li> <li>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i></li> </ul>

## traceoptions

<b>Syntax</b>	<pre> traceoptions {   file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt;;   flag all;   flag certificates;   flag database;   flag general;   flag ike;   flag parse;   flag policy-manager;   flag routing-socket;   flag timer;   level   no-remote-trace } </pre>
<b>Hierarchy Level</b>	<p>[edit security], [edit services ipsec-vpn]</p> <p>Trace options can be configured at either the [edit security] or the [edit services ipsec-vpn] hierarchy level, but not at both levels.</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure security trace options.</p> <p>To specify more than one trace option, include multiple <b>flag</b> statements. Trace option output is recorded in the <code>/var/log/kmd</code> file.</p>
<div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> The <code>traceoptions</code> statement is not supported on QFabric systems.</p> </div> </div>	
<b>Options</b>	<p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file (for example, <code>kmd</code>) reaches its maximum size, it is renamed <code>kmd.0</code>, then <code>kmd.1</code>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the <b>size</b> option.</p> <p><b>Range:</b> 2 through 1000 files</p> <p><b>Default:</b> 0 files</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each trace file, in kilobytes (KB). When a trace file (for example, <code>kmd</code>) reaches this size, it is renamed, <code>kmd.0</code>, then <code>kmd.1</code> and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Default:</b> 1024 KB</p>

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

- **all**—Trace all security events.
- **certificates**—Trace certificate events.
- **database**—Trace database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.
- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **timer**—Trace internal timer events.

**level *level***—(Optional) Set traceoptions level.

- **all**—match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match conditions that should be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

**no-remote-trace**—(Optional) Disable remote tracing

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

<b>Related Documentation</b>	• <i>Configuring Tracing Operations for Security Services</i>
------------------------------	---------------------------------------------------------------



## CHAPTER 6

# Operational Command (Port Mirroring)

- `show analyzer`

## show analyzer

<b>Syntax</b>	<b>show analyzer</b> < <i>analyzer-name</i> >
<b>Release Information</b>	Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	Display information about port mirroring.
<b>Options</b>	<i>analyzer-name</i> —(Optional) Displays the status of a specific analyzer (port-mirroring configuration).
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Understanding Layer 2 Port Mirroring</i></li> <li><a href="#">Port Mirroring Constraints and Limitations on page 6</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show analyzer on page 96</a>
<b>Output Fields</b>	<a href="#">Table 4 on page 96</a> describes the output fields for the <b>show analyzer</b> command. Output fields are listed in the approximate order in which they appear.

**Table 4: show analyzer Output Fields**

Field Name	Field Description
Analyzer name	Name of the analyzer.
Output interface	Local interface to which mirror packets are sent. If you configure an output interface, you cannot also configure an output VLAN.
Output VLAN	VLAN to which mirror packets are sent. If you configure an output VLAN, you cannot also configure an output interface.
Egress monitored interfaces	Interfaces for which egress traffic is mirrored.
Egress monitored VLANs	VLANs for which egress traffic is mirrored.
Ingress monitored interfaces	Interfaces for which ingress traffic is mirrored.
Ingress monitored VLANs	VLANs for which ingress traffic is mirrored.

## Sample Output

### show analyzer

```

user@switch> show analyzer
Analyzer name           : employee-monitor
Output interface        : ge-0/0/10.0
Output VLAN             : remote-analyzer
Egress monitored interfaces : ge-0/0/7.0

```



```
Ingress monitored interfaces : ge-0/0/8.0  
Ingress monitored interfaces : ge-0/0/9.0
```



## CHAPTER 7

# Operational Commands (DHCP Local Server)

- clear dhcp server binding
- clear dhcp server statistics
- clear dhcpv6 server binding
- clear dhcpv6 server statistics
- request dhcp server reconfigure
- request dhcpv6 server reconfigure
- request system reboot
- show dhcp server binding
- show dhcp server statistics
- show dhcpv6 server binding
- show dhcpv6 server statistics

## clear dhcp server binding

**Syntax** `clear dhcp server binding`  
`<address>`  
`<all>`  
`<interface interface-name>`  
`<interfaces-vlan>`  
`<interfaces-wildcard>`  
`<logical-system logical-system-name>`  
`<routing-instance routing-instance-name>`

**Release Information** Command introduced in Junos OS Release 9.0.  
Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

**Description** Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table on the extended DHCP local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

**Options** *address*—(Optional) Clear the binding state for the DHCP client, using one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

*all*—(Optional) Clear the binding state for all DHCP clients.

*interface interface-name*—(Optional) Clear the binding state for DHCP clients on the specified interface.



**NOTE:** This option clears all bindings whose initial login requests were received over the specified interface. Dynamic demux login requests are not received over the dynamic demux interface, but rather the underlying interface of the dynamic demux interface. To clear a specific dynamic demux interface, use the *ip-address* or *mac-address* options.

*interfaces-vlan*—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.

*interfaces-wildcard*—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (\*).

**logical-system** *logical-system-name*—(Optional) Clear the binding state for DHCP clients on the specified logical system.

**routing-instance** *routing-instance-name*—(Optional) Clear the binding state for DHCP clients on the specified routing instance.

**Required Privilege Level** view

**Related Documentation**

- [Clearing DHCP Bindings for Subscriber Access](#)
- [show dhcp server binding on page 117](#)

**List of Sample Output**

- [clear dhcp server binding <ip-address> on page 101](#)
- [clear dhcp server binding all on page 101](#)
- [clear dhcp server binding interface on page 102](#)
- [clear dhcp server binding <interfaces-vlan> on page 102](#)
- [clear dhcp server binding <interfaces-wildcard> on page 102](#)

**Output Fields** See [show dhcp server binding](#) for an explanation of output fields.

## Sample Output

### clear dhcp server binding <ip-address>

The following sample output displays the address bindings in the DHCP client table on the extended DHCP local server before and after the **clear dhcp server binding** command is issued.

```
user@host> show dhcp server binding
```

```
2 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.1	90:00:00:01:00:01	active	2007-01-17 11:38:47 PST
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

```
user@host> clear dhcp server binding 10.20.32.1
```

```
user@host> show dhcp server binding
```

```
1 clients, (0 bound, 0 selecting, 0 renewing, 0 rebinding)
```

IP address	Hardware address	Type	Lease expires at
100.20.32.3	90:00:00:02:00:01	active	2007-01-17 11:38:41 PST

### clear dhcp server binding all

The following command clears all DHCP local server bindings:

```
user@host> clear dhcp server binding all
```

### clear dhcp server binding interface

The following command clears DHCP local server bindings on a specific interface:

```
user@host> clear dhcp server binding interface fe-0/0/2
```

### clear dhcp server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP local server bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp server binding ae0
```

### clear dhcp server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP local server bindings over a specific interface:

```
user@host> clear dhcp server binding ge-1/0/0.*
```

## clear dhcp server statistics

<b>Syntax</b>	<code>clear dhcp server statistics</code> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	Clear all extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Clear the statistics for DHCP clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Clear the statistics for DHCP clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear dhcp server statistics on page 103</a>
<b>Output Fields</b>	See <a href="#">show dhcp server statistics</a> for an explanation of output fields.

## Sample Output

### clear dhcp server statistics

The following sample output displays the extended DHCP local server statistics before and after the **clear dhcp server statistics** command is issued.

```

user@host> show dhcp server statistics
Packets dropped:
    Total                1
    Lease Time Violation 1

Messages received:
    BOOTREQUEST          89163
    DHCPDECLINE           0
    DHCPDISCOVER          8110
    DHCPINFORM            0
    DHCPRELEASE           0
    DHCPREQUEST          81053

Messages sent:
    BOOTREPLY             32420
    DHCPOFFER             8110
    DHCPACK               8110
    DHCPNAK               8100

user@host> clear dhcp server statistics
user@host> show dhcp server statistics

```

Packets dropped:	
Total	0
Messages received:	
BOOTREQUEST	0
DHCPDECLINE	0
DHCPDISCOVER	0
DHCPINFORM	0
DHCPRELEASE	0
DHCPREQUEST	0
Messages sent:	
BOOTREPLY	0
DHCPOFFER	0
DHCPACK	0
DHCPNAK	0



## clear dhcpv6 server binding

<b>Syntax</b>	<pre>clear dhcpv6 server binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.6.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p>
<b>Description</b>	Clear the binding state of a Dynamic Host Configuration Protocol for IPv6 (DHCPv6) client from the client table on the extended DHCPv6 local server.
<b>Options</b>	<p><b>address</b>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) Clear bindings on a set of interfaces. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Clearing DHCP Bindings for Subscriber Access</a></li> <li>• <a href="#">show dhcpv6 server binding on page 126</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">clear dhcpv6 server binding all on page 106</a></p> <p><a href="#">clear dhcpv6 server binding &lt;ipv6-prefix&gt; on page 106</a></p>

[clear dhcpv6 server binding interface on page 106](#)  
[clear dhcpv6 server binding <interfaces-vlan> on page 106](#)  
[clear dhcpv6 server binding <interfaces-wildcard> on page 106](#)

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

## Sample Output

### clear dhcpv6 server binding all

The following command clears all DHCPv6 local server bindings:

```
user@host> clear dhcpv6 server binding all
```

### clear dhcpv6 server binding <ipv6-prefix>

The following command clears DHCPv6 local server bindings for a specific IPv6 prefix:

```
user@host> clear dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
```

### clear dhcpv6 server binding interface

The following command clears DHCPv6 local server bindings on a specific interface:

```
user@host> clear dhcpv6 server binding interface fe-0/0/2
```

### clear dhcpv6 server binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 local server bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 server binding interface ae0
```

### clear dhcpv6 server binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 local server bindings over a specific interface:

```
user@host> clear dhcpv6 server binding ge-1/0/0.*
```

## clear dhcpv6 server statistics

---


<b>Syntax</b>	<pre>clear dhcpv6 server statistics &lt;interface <i>interface-name</i>&gt; &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 9.6.
<b>Description</b>	Clear all extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Clear the statistics for DHCPv6 clients on the specified logical system. If you do not specify a logical system, statistics are cleared for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Clear the statistics for DHCPv6 clients on the specified routing instance. If you do not specify a routing instance, statistics are cleared for the default routing instance.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show dhcpv6 server statistics on page 132</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear dhcpv6 server statistics on page 107</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear dhcpv6 server statistics

```
user@host> clear dhcpv6 server statistics
```

## request dhcp server reconfigure

<b>Syntax</b>	<code>request dhcp server reconfigure (all   <i>address</i>   interface <i>interface-name</i>   logical-system <i>logical-system-name</i>   routing-instance <i>routing-instance-name</i>)</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.0. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Initiate reconfiguration processing for the specified DHCP clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the <b>clear dhcp server binding</b> command.</p> <p>When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfiguring state and the local server sends a forcerenew message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the forcerenew message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the <b>clear-on-abort</b> statement to configure the client to be cleared when reconfiguration fails.</p>
<b>Options</b>	<p><b>all</b>—Initiate reconfiguration for all DHCP clients.</p> <p><b><i>address</i></b>—Initiate reconfiguration for DHCP client with the specified IP address or MAC address.</p> <p><b>interface <i>interface-name</i></b>—Initiate reconfiguration for all DHCP clients on this logical interface (clients whose initial login requests were received over the specified interface).</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> <b>NOTE:</b> You cannot use the <b>interface <i>interface-name</i></b> option with the <b>request dhcp server reconfigure</b> command for DHCP passive clients (clients that are added as a result of DHCP snooped packets). For passive clients, the interface is not guaranteed to be the next-hop interface to the client, as is the case for active clients.</p> </div> <p><b>logical-system <i>logical-system-name</i></b>—Initiate reconfiguration for all DHCP clients on the specified logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—Initiate reconfiguration reconfigured for all DHCP clients in the specified routing instance.</p>
<b>Required Privilege Level</b>	view

**Related Documentation** • *Configuring Dynamic Client Reconfiguration of Extended Local Server Clients*

**List of Sample Output** [request dhcp server reconfigure on page 109](#)

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

## Sample Output

[request dhcp server reconfigure](#)

```
user@host> request dhcp server reconfigure interface fe-0/0/0.100
```

## request dhcpv6 server reconfigure

<b>Syntax</b>	<b>request dhcpv6 server reconfigure</b> ( <b>all</b>   <b>address</b>   <b>client-id</b>   <b>interface</b> <i>interface-name</i>   <b>logical-system</b> <i>logical-system-name</i>   <b>routing-instance</b> <i>routing-instance-name</i>   <b>session-id</b> )
<b>Release Information</b>	Command introduced in Junos OS Release 10.4. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	<p>Initiate reconfiguration processing for the specified DHCPv6 clients if they are in the bound state. If the clients are in the reconfiguring state, this command has no effect. If the clients are in any state other than bound or reconfiguring, this command has the same effect as the <b>clear dhcpv6 server binding</b> command.</p> <p>When the local server state machine starts the reconfiguration process on a bound client, the client transitions to the reconfigure state and the local server sends a reconfigure message to the client. Because the client was in the bound state before entering the reconfiguring state, all subscriber (or DHCP client) services, such as forwarding and statistics, continue to work. An exponential back-off timer determines the interval at which the reconfigure message is sent. If the final attempt is unsuccessful, the client is returned to its original state by default. You can optionally include the <b>clear-on-abort</b> statement to configure the client to be cleared when reconfiguration fails.</p>
<b>Options</b>	<p><b>all</b>—Initiate reconfiguration for all DHCPv6 clients.</p> <p><b>address</b>—Initiate reconfiguration for DHCPv6 client with the specified IPv6 address.</p> <p><b>client-id</b>—Initiate reconfiguration for DHCPv6 client with the specified client ID.</p> <p><b>interface</b> <i>interface-name</i>—Initiate reconfiguration for all DHCPv6 clients on this logical interface (clients whose initial login requests were received over the specified interface).</p> <p><b>logical-system</b> <i>logical-system-name</i>—Initiate reconfiguration for all DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance</b> <i>routing-instance-name</i>—Initiate reconfiguration reconfigured for all DHCPv6 clients in the specified routing instance.</p> <p><b>session-id</b>—Initiate reconfiguration for DHCPv6 client with the specified session ID.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Configuring Dynamic Client Reconfiguration of Extended Local Server Clients</i></li> </ul>
<b>List of Sample Output</b>	<a href="#">request dhcpv6 server reconfigure on page 111</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

request dhcpv6 server reconfigure

```
user@host> request dhcpv6 server reconfigure 2001::2/16
```

## request system reboot

---

<b>List of Syntax</b>	<a href="#">Syntax on page 112</a> <a href="#">Syntax (EX Series Switches) on page 112</a> <a href="#">Syntax (TX Matrix Router) on page 112</a> <a href="#">Syntax (TX Matrix Plus Router) on page 112</a> <a href="#">Syntax (MX Series Router) on page 112</a>
<b>Syntax</b>	<code>request system reboot</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk   removable-compact-flash   usb)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code>
<b>Syntax (EX Series Switches)</b>	<code>request system reboot</code> <code>&lt;all-members&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;local&gt;</code> <code>&lt;media (external   internal)&gt;</code> <code>&lt;member <i>member-id</i>&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;slice <i>slice</i>&gt;</code>
<b>Syntax (TX Matrix Router)</b>	<code>request system reboot</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   scc&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code>
<b>Syntax (TX Matrix Plus Router)</b>	<code>request system reboot</code> <code>&lt;all-chassis   all-lcc   lcc <i>number</i>   sfc <i>number</i>&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;media (compact-flash   disk)&gt;</code> <code>&lt;message "<i>text</i>"&gt;</code> <code>&lt;other-routing-engine&gt;</code> <code>&lt;partition (1   2   alternate)&gt;</code>
<b>Syntax (MX Series Router)</b>	<code>request system reboot</code> <code>&lt;all-members&gt;</code> <code>&lt;at <i>time</i>&gt;</code> <code>&lt;both-routing-engines&gt;</code> <code>&lt;in <i>minutes</i>&gt;</code> <code>&lt;local&gt;</code>



```

<media (external | internal)>
<member member-id>
<message "text">
<other-routing-engine>

```

<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option <b>other-routing-engine</b> introduced in Junos OS Release 8.0.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option <b>sfc</b> introduced for the TX Matrix Plus router in Junos OS Release 9.6.</p> <p>Option <b>both-routing-engines</b> introduced in Junos OS Release 12.1.</p>
<b>Description</b>	Reboot the software.
<b>Options</b>	<p><b>none</b>—Reboot the software immediately.</p> <p><b>all-chassis</b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, reboot all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, reboot all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.</p> <p><b>all-lcc</b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, reboot all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, reboot all T1600 routers (or line-card chassis) connected to the TX Matrix Plus router.</p> <p><b>all-members</b>—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on all members of the Virtual Chassis configuration.</p> <p><b>at <i>time</i></b>—(Optional) Time at which to reboot the software, specified in one of the following ways:</p> <ul style="list-style-type: none"> <li>• <b>now</b>—Stop or reboot the software immediately. This is the default.</li> <li>• <b>+<i>minutes</i></b>—Number of minutes from now to reboot the software.</li> <li>• <b><i>yyymmddhhmm</i></b>—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.</li> <li>• <b><i>hh:mm</i></b>—Absolute time on the current day at which to stop the software, specified in 24-hour time.</li> </ul> <p><b>both-routing-engines</b>—(Optional) Reboot both Routing Engines at the same time.</p> <p><b>in <i>minutes</i></b>—(Optional) Number of minutes from now to reboot the software. This option is an alias for the <b>at +<i>minutes</i></b> option.</p> <p><b>lcc <i>number</i></b>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router or TX Matrix Plus router, the number of a specified line card chassis connected to the TX Matrix or TX Matrix Plus router, respectively. Replace <b><i>number</i></b> with a value from 0 through 3.</p> <p><b>local</b>—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the local Virtual Chassis member.</p>

**media (compact-flash | disk)**—(Optional) Boot medium for next boot.

**media (external | internal)**—(EX Series switches and MX Series routers only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.
- **internal**—Reboot the internal flash device.

**member *member-id***—(EX4200 switches and MX Series routers only) (Optional) Reboot the software on the specified member of the Virtual Chassis configuration. For EX4200 switches, replace ***member-id*** with a value from 0 through 9. For an MX Series Virtual Chassis, replace ***member-id*** with a value of 0 or 1.

**message "*text*"**—(Optional) Message to display to all system users before stopping or rebooting the software.

**other-routing-engine**—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

**partition**—(TX Matrix Plus routers only) (Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition.

**scc**—(TX Matrix routers only) (Optional) Reboot the Routing Engine on the TX Matrix switch-card chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted.

**sfc *number***—(TX Matrix Plus routers only) (Optional) Reboot the Routing Engine on the TX Matrix Plus switch-fabric chassis. If you issue the command from re0, re0 is rebooted. If you issue the command from re1, re1 is rebooted. Replace ***number*** with 0.

**slice *slice***—(EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

**Additional Information** Reboot requests are recorded in the system log files, which you can view with the **show log** command (see *show log*). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see *show system processes*).

On a TX Matrix or TX Matrix Plus router, if you issue the **request system reboot** command on the master Routing Engine, all the master Routing Engines connected to the routing matrix are rebooted. If you issue this command on the backup Routing Engine, all the backup Routing Engines connected to the routing matrix are rebooted.



**NOTE:** Before issuing the **request system reboot** command on a TX Matrix Plus router with no options or the **all-chassis**, **all-lcc**, **lcc number**, or **sfc** options, verify that master Routing Engine for all routers in the routing matrix are in the same slot number. If the master Routing Engine for a line-card chassis is in a different slot number than the master Routing Engine for a TX Matrix Plus router, the line-card chassis might become logically disconnected from the routing matrix after the **request system reboot** command.



**NOTE:** To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

**Required Privilege Level** maintenance

**Related Documentation**

- *clear system reboot*
- *request system halt*
- *request system reboot*
- *Rebooting and Halting a Device*

**List of Sample Output**

- [request system reboot on page 115](#)
- [request system reboot \(at 2300\) on page 115](#)
- [request system reboot \(in 2 Hours\) on page 116](#)
- [request system reboot \(Immediately\) on page 116](#)
- [request system reboot \(at 1:20 AM\) on page 116](#)

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

## Sample Output

### request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

### request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes
```

```
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

#### request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

#### request system reboot (Immediately)

```
user@host> request system reboot at now
```

#### request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

## show dhcp server binding

**Syntax**    `show dhcp server binding`  
               `<address>`  
               `<interfaces-vlan><brief | detail | summary>`  
               `<interface interface-name>`  
               `<interfaces-vlan>`  
               `<interfaces-wildcard>`  
               `<logical-system logical-system-name>`  
               `<routing-instance routing-instance-name>`

**Release Information**    Command introduced in Junos OS Release 9.0.  
                               Options *interfaces-vlan* and *interfaces-wildcard* added in Junos OS Release 12.1.

**Description**    Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol (DHCP) local server.



**NOTE:** If you delete the DHCP server configuration, DHCP server bindings might still remain. To ensure that DHCP bindings are removed, issue the `clear dhcp server binding` command before you delete the DHCP server configuration.

**Options**    **address**—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

**brief | detail | summary**—(Optional) Display the specified level of output about active client bindings. The default is **brief**, which produces the same output as `show dhcp server binding`.

**interface interface-name**—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

**interfaces-vlan**—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

**interfaces-wildcard**—(Optional) The set of interfaces on which to show the binding state information. This option supports the use of the wildcard character (\*).

**logical-system logical-system-name**—(Optional) Display information about active client bindings for DHCP clients on the specified logical system.

**routing-instance routing-instance-name**—(Optional) Display information about active client bindings for DHCP clients on the specified routing instance.

**Required Privilege Level** view

**Related Documentation**

- *Clearing DHCP Bindings for Subscriber Access*
- *Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration*
- [clear dhcp server binding on page 100](#)

**List of Sample Output**

[show dhcp server binding on page 120](#)  
[show dhcp server binding detail on page 121](#)  
[show dhcp server binding detail \(ACI Interface Set Configured\) on page 121](#)  
[show dhcp server binding interface <vlan-id> on page 122](#)  
[show dhcp server binding interface <svlan-id> on page 122](#)  
[show dhcp server binding <ip-address> on page 122](#)  
[show dhcp server binding <session-id> on page 122](#)  
[show dhcp server binding summary on page 122](#)  
[show dhcp server binding <interfaces-vlan> on page 122](#)  
[show dhcp server binding <interfaces-wildcard> on page 122](#)

**Output Fields** [Table 5 on page 118](#) lists the output fields for the **show dhcp server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 5: show dhcp server binding Output Fields**

Field Name	Field Description	Level of Output
<i>number</i> clients, ( <i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	<b>summary</b>
IP address	IP address of the DHCP client.	<b>brief detail</b>
Session Id	Session ID of the subscriber session.	<b>brief detail</b>
Hardware address	Hardware address of the DHCP client.	<b>brief detail</b>
Expires	Number of seconds in which lease expires.	<b>brief detail</b>

Table 5: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the address binding table on the extended DHCP local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>FORCERENEW</b>—Client has received forcerenew message from server.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCP server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCP servers.</li> </ul>	<b>brief</b> <b>detail</b>
<b>Interface</b>	Interface on which the request was received.	<b>brief</b>
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>
<b>Lease Expires in</b>	Number of seconds in which lease expires.	<b>detail</b>
<b>Lease Start</b>	Date and time at which the client's IP address lease started.	<b>detail</b>
<b>Lease time violated</b>	Lease time violation has occurred.	<b>detail</b>
<b>Last Packet Received</b>	Date and time at which the router received the last packet.	<b>detail</b>
<b>Incoming Client Interface</b>	Client's incoming interface.	<b>detail</b>
<b>Client Interface Svlan Id</b>	S-VLAN ID of the client's incoming interface.	<b>detail</b>
<b>Client Interface Vlan Id</b>	VLAN ID of the client's incoming interface.	<b>detail</b>
<b>Demux Interface</b>	Name of the IP demultiplexing (demux) interface.	<b>detail</b>
<b>Server IP Address or Server Identifier</b>	IP address of DHCP server.	<b>detail</b>
<b>Server Interface</b>	Interface of DHCP server.	<b>detail</b>
<b>Client Pool Name</b>	Name of address pool used to assign client IP address lease.	<b>detail</b>

Table 5: show dhcp server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Liveness Detection State	<p>State of the liveness detection status for a subscriber's Bidirectional Forwarding Detection (BFD) protocol session:</p> <p><b>NOTE:</b> This output field displays status only when liveness detection has been explicitly configured for a subscriber and the liveness detection protocol is actively functioning for that subscriber.</p> <ul style="list-style-type: none"> <li><b>DOWN</b>—Liveness detection has been enabled for a subscriber but the broadband network gateway (BNG) detects that the liveness detection session for the BFD protocol is in the <b>DOWN</b> state. A liveness detection session that was previously in an <b>UP</b> state has transitioned to a <b>DOWN</b> state, beginning with a liveness detection failure, and ending with the deletion of the client binding. The <b>DOWN</b> state is reported only during this transition period of time.</li> <li><b>UNKNOWN</b>—Liveness detection has been enabled for a subscriber but the actual liveness detection state has not yet been determined. The <b>UNKNOWN</b> state is reported after a DHCP subscriber initially logs in while the underlying liveness detection protocol handshake, such as BFD, is still processing and the BFD session has not yet reached the <b>UP</b> state.</li> <li><b>UP</b>—Liveness detection has been enabled for a subscriber, and the BNG and the subscriber or client have <i>both</i> determined that the liveness detection session for the BFD protocol is in the <b>UP</b> state.</li> <li><b>WENT_DOWN</b>—State is functionally equivalent to the <b>DOWN</b> state. A liveness detection session that was previously in an <b>UP</b> state has transitioned to a <b>DOWN</b> state implying a liveness detection failure. The <b>WENT_DOWN</b> state applies to the internal distribution of the liveness detection mechanism between the Junos DHCP Daemon for Subscriber Services (JDHCPd), the BFD plug-in within the Broadband Edge Subscriber Management Daemon (BBE-SMGD), and the Packet Forwarding Engine.</li> </ul>	detail
ACI Interface Set Name	Internally generated name of the dynamic agent circuit identifier (ACI) interface set.	detail
ACI Interface Set Index	Index number of the dynamic ACI interface set.	detail
ACI Interface Set Session ID	Identifier of the dynamic ACI interface set entry in the session database.	detail

## Sample Output

### show dhcp server binding

```
user@host> show dhcp server binding
```



IP address	Session Id	Hardware address	Expires	State	Interface
100.20.20.15	6	00:10:94:00:00:01	86180	BOUND	ge-1/0/0.0
100.20.20.16	7	00:10:94:00:00:02	86180	BOUND	ge-1/0/0.0
100.20.20.17	8	00:10:94:00:00:03	86180	BOUND	ge-1/0/0.0
100.20.20.18	9	00:10:94:00:00:04	86180	BOUND	ge-1/0/0.0
100.20.20.19	10	00:10:94:00:00:05	86180	BOUND	ge-1/0/0.0

### show dhcp server binding detail

```

user@host> show dhcp server binding detail
Client IP Address: 100.20.20.15
    Hardware Address:      00:10:94:00:00:01
    State:                 BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

    Lease Expires:        2009-07-21 10:10:25 PDT
    Lease Expires in:     86151 seconds
    Lease Start:          2009-07-20 10:10:25 PDT
    Incoming Client Interface: ge-1/0/0.0
    Server Ip Address:     100.20.20.9
    Server Interface:      none
    Session Id:           6
    Client Pool Name:      6
    Liveness Detection State: UP
Client IP Address: 100.20.20.16
    Hardware Address:      00:10:94:00:00:02
    State:                 BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)

    Lease Expires:        2009-07-21 10:10:25 PDT
    Lease Expires in:     86151 seconds
    Lease Start:          2009-07-20 10:10:25 PDT
    Lease time violated:   yes
    Incoming Client Interface: ge-1/0/0.0
    Server Ip Address:     100.20.20.9
    Server Interface:      none
    Session Id:           7
    Client Pool Name:      7
    Liveness Detection State: UP

```

### show dhcp server binding detail (ACI Interface Set Configured)

```

user@host> show dhcp server binding detail
Client IP Address: 100.20.22.14
    Hardware Address:      00:00:64:34:01:02
    State:                 BOUND(LOCAL_SERVER_STATE_BOUND)
    Lease Expires:        2012-03-13 09:53:32 PDT
    Lease Expires in:     82660 seconds
    Lease Start:          2012-03-12 10:23:32 PDT
    Last Packet Received: 2012-03-12 10:23:32 PDT
    Incoming Client Interface: demux0.1073741827
    Client Interface Svlan Id: 1802
    Client Interface Vlan Id: 302
    Demux Interface:       demux0.1073741832
    Server Identifier:     100.20.200.202
    Session Id:           11
    Client Pool Name:      poolA
    Client Profile Name:   DEMUXprofile
    Liveness Detection State: UP

```

```

ACI Interface Set Name:      aci-1002-demux0.1073741827
ACI Interface Set Index:    2
ACI Interface Set Session ID: 6

```

#### show dhcp server binding interface <vlan-id>

```

user@host> show dhcp server binding interface ge-1/1/0:100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01 86124    BOUND  ge-1/1/0:100

```

#### show dhcp server binding interface <svlan-id>

```

user@host> show dhcp server binding interface ge-1/1/0:10-100
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.16    7          00:10:94:00:00:02 86124    BOUND  ge-1/1/0:10-100

```

#### show dhcp server binding <ip-address>

```

user@host> show dhcp server binding 100.20.20.19
IP address      Session Id  Hardware address  Expires  State  Interface
100.20.20.19    10         00:10:94:00:00:05 86081    BOUND  ge-1/0/0.0

```

#### show dhcp server binding <session-id>

```

user@host> show dhcp server binding 6
IP address      Session Id  Hardware address  Expires  State  Interface
200.20.20.15    6          00:10:94:00:00:01 86124    BOUND  ge-1/0/0.0

```

#### show dhcp server binding summary

```

user@host> show dhcp server binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```

#### show dhcp server binding <interfaces-vlan>

```

user@host> show dhcp server binding ge-1/0/0:100-200
IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.17    42         00:10:94:00:00:02 86346    BOUND  ge-1/0/0.1073741827
192.168.0.16    41         00:10:94:00:00:01 86346    BOUND  ge-1/0/0.1073741827

```

#### show dhcp server binding <interfaces-wildcard>

```

user@host> show dhcp server binding ge-1/3/*
IP address      Session Id  Hardware address  Expires  State  Interface
192.168.0.9     24         00:10:94:00:00:04 86361    BOUND  ge-1/3/0.110
192.168.0.8     23         00:10:94:00:00:03 86361    BOUND  ge-1/3/0.110
192.168.0.7     22         00:10:94:00:00:02 86361    BOUND  ge-1/3/0.110

```

## show dhcp server statistics

---

<b>Syntax</b>	<b>show dhcp server statistics</b> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.0.
<b>Description</b>	Display extended Dynamic Host Configuration Protocol (DHCP) local server statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display information about extended DHCP local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display information about extended DHCP local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear dhcp server statistics on page 103</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show dhcp server statistics on page 124</a>
<b>Output Fields</b>	<a href="#">Table 6 on page 124</a> lists the output fields for the <b>show dhcp server statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 6: show dhcp server statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCP local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP local server</li> <li>• <b>Authentication</b>—Number of packets discarded because they could not be authenticated</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Dynamic profile</b>—Number of packets discarded due to dynamic profile information</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP local server could not send</li> </ul>
<b>Messages received</b>	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHCPNACK</b>—Number of DHCP NACK PDUs transmitted</li> <li>• <b>DHCPFORCERENEW</b>—Number of DHCP FORCERENEW PDUs transmitted</li> </ul>

## Sample Output

### show dhcp server statistics

```

user@host> show dhcp server statistics
Packets dropped:
    Total                  1

```

Lease Time Violation	1
Messages received:	
BOOTREQUEST	25
DHCPDECLINE	0
DHCPDISCOVER	10
DHCPINFORM	0
DHCPRELEASE	4
DHCPREQUEST	10
Messages sent:	
BOOTREPLY	20
DHCPOFFER	10
DHCPACK	10
DHCPNAK	0
DHCPFORCERENEW	0

## show dhcpv6 server binding

<b>Syntax</b>	<pre>show dhcpv6 server binding &lt;address&gt; &lt;brief   detail   summary&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 9.6.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p>
<b>Description</b>	Display the address bindings in the client table on the extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server.
<b>Options</b>	<p><b>address</b>—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>brief   detail   summary</b>—(Optional) Display the specified level of output about active client bindings. The default is <b>brief</b>, which produces the same output as <b>show dhcpv6 server binding</b>.</p> <p><b>interface interface-name</b>—(Optional) Display information about active client bindings on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p><b>interfaces-wildcard</b>—(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Display information about active client bindings for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Display information about active client bindings for DHCPv6 clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li> <li>• <a href="#">clear dhcpv6 server binding on page 105</a></li> </ul>

**List of Sample Output**

- [show dhcpv6 server binding on page 128](#)
- [show dhcpv6 server binding detail on page 128](#)
- [show dhcpv6 server binding interface on page 129](#)
- [show dhcpv6 server binding interface detail on page 129](#)
- [show dhcpv6 server binding \(IPv6 Prefix\) on page 130](#)
- [show dhcpv6 server binding \(Session ID\) on page 130](#)
- [show dhcpv6 server binding \(Interfaces VLAN\) on page 130](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 130](#)
- [show dhcpv6 server binding \(Interfaces Wildcard\) on page 130](#)
- [show dhcpv6 server binding summary on page 131](#)

**Output Fields** Table 7 on page 127 lists the output fields for the **show dhcpv6 server binding** command. Output fields are listed in the approximate order in which they appear.

**Table 7: show dhcpv6 server binding Output Fields**

Field Name	Field Description	Level of Output
<i>number clients</i> , ( <i>number init</i> , <i>number bound</i> , <i>number selecting</i> , <i>number requesting</i> , <i>number renewing</i> , <i>number releasing</i> )	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	<b>summary</b>
<b>Prefix</b>	Client's DHCPv6 prefix, or prefix used to support multiple address assignment.	<b>brief detail</b>
<b>Session Id</b>	Session ID of the subscriber session.	<b>brief detail</b>
<b>Expires</b>	Number of seconds in which lease expires.	<b>brief detail</b>
<b>State</b>	State of the address binding table on the extended DHCPv6 local server: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>RECONFIGURE</b>—Server has sent reconfigure message to client.</li> <li>• <b>RELEASE</b>—Client is releasing IP address lease.</li> <li>• <b>RENEWING</b>—Client sending request to renew IP address lease.</li> <li>• <b>REQUESTING</b>—Client requesting a DHCPv6 server.</li> <li>• <b>SELECTING</b>—Client receiving offers from DHCPv6 servers.</li> </ul>	<b>brief detail</b>
<b>Interface</b>	Interface on which the DHCPv6 request was received.	<b>brief</b>
<b>Client IPv6 Address</b>	Client's IPv6 address.	<b>detail</b>
<b>Client IPv6 Prefix</b>	Client's IPv6 prefix.	<b>detail</b>
<b>Client DUID</b>	Client's DHCP Unique Identifier (DUID).	<b>brief detail</b>
<b>Lease expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>

Table 7: show dhcpv6 server binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease expires in	Number of seconds in which lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which client's IPv6 prefix expires.	detail
Lease Start	Date and time at which the client's address lease was obtained.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server IP Address	IP address of DHCPv6 server.	detail
Server Interface	Interface of DHCPv6 server.	detail
Client Pool Name	Address pool used to assign IPv6 address.	detail
Client Prefix Pool Name	Address pool used to assign IPv6 prefix.	detail
Client Id length	Length of the DHCPv6 client ID, in bytes.	detail
Client Id	ID of the DHCPv6 client.	detail

## Sample Output

### show dhcpv6 server binding

```

user@host> show dhcpv6 server binding
Prefix          Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 6 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:01
2001:bd8:1111:2222::/64 7 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:02
2001:bd8:1111:2222::/64 8 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03
2001:bd8:1111:2222::/64 9 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:04
2001:bd8:1111:2222::/64 10 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:05
2002::1/74 11 86321 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c1-00:10:94:00:00:06

```

### show dhcpv6 server binding detail

```

user@host> show dhcpv6 server binding detail

```



```

Session Id: 6
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:01

  State:
BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
  Lease Expires:           2009-07-21 10:41:15 PDT
  Lease Expires in:        86308 seconds
  Preferred Lease Expires: 2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:             2009-07-20 10:41:15 PDT
  Lease time violated:     yes
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:       0.0.0.0
  Server Interface:        none
  Client Id Length:        14
  Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0001

```

```

Session Id: 7
  Client IPv6 Address:     2002::1/128
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

  State:
BOUND(LOCAL_SERVER_STATE_BOUND_ON_INTF_DELETE)
  Lease Expires:           2009-07-21 10:41:15 PDT
  Lease Expires in:        86308 seconds
  Preferred Lease Expires: 2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:             2009-07-20 10:41:15 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:       0.0.0.0
  Client Pool Name:        bos-v6-pool
  Client Prefix Pool Name: bos-v6-prefix-pool
  Client Id Length:        14
  Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding interface

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101
Prefix      Session Id Expires State Interface Client DUID
2001:bd8:1111:2222::/64 1      86055   BOUND   ge-1/0/0.100
LL_TIME0x1-0x4b0a53b9-00:10:94:00:00:01

```

### show dhcpv6 server binding interface detail

```

user@host> show dhcpv6 server binding interface ge-1/0/0:10-101 detail
Session Id: 7
  Client IPv6 Prefix:      2001:bd8:1111:2222::/64
  Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

  State:                   BOUND(bound)
  Lease Expires:           2009-07-21 10:41:15 PDT
  Lease Expires in:        86136 seconds
  Preferred Lease Expires: 2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:             2009-07-20 10:41:15 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address:       0.0.0.0

```

```

Server Interface:          none
Client Id Length:         14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding (IPv6 Prefix)

```

user@host> show dhcpv6 server binding 14/0x00010001/0x02b3be8f/0x00109400/0x0005
detail
Session Id: 7
Client IPv6 Prefix:      2001:bd8:1111:2222::/64
Client DUID:             LL_TIME0x1-0x2e159c0-00:10:94:00:00:02

State:                   BOUND(bound)
Lease Expires:           2009-07-21 10:41:15 PDT
Lease Expires in:        86136 seconds
Preferred Lease Expires: 2012-07-24 00:18:14 UTC
Preferred Lease Expires in: 600 seconds
Lease Start:             2009-07-20 10:41:15 PDT
Incoming Client Interface: ge-1/0/0.0
Server Ip Address:       0.0.0.0
Server Interface:        none
Client Id Length:        14
Client Id:
/0x00010001/0x02e159c0/0x00109400/0x0002

```

### show dhcpv6 server binding (Session ID)

```

user@host> show dhcpv6 server binding 8
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 8      86235 BOUND ge-1/0/0.0
LL_TIME0x1-0x2e159c0-00:10:94:00:00:03

```

### show dhcpv6 server binding (Interfaces VLAN)

```

user@host> show dhcpv6 server binding ge-1/0/0:100-200
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 11      87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 12      87583 BOUND ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

### show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding demux0
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 30      79681 BOUND demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 31      79681 BOUND demux0.1073741825
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32 32      79681 BOUND demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

### show dhcpv6 server binding (Interfaces Wildcard)

```

user@host> show dhcpv6 server binding ge-1/3/*
Prefix      Session Id Expires State Interface Client DUID
2001:DB8::/32 22      79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32 33      79681 BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

```
2001:CB9::/32      24      79681    BOUND    ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
```

#### show dhcpv6 server binding summary

```
user@host> show dhcpv6 server binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)
```

## show dhcpv6 server statistics

---

<b>Syntax</b>	<b>show dhcpv6 server statistics</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 9.6.
<b>Description</b>	Display extended Dynamic Host Configuration Protocol for IPv6 (DHCPv6) local server statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Display information about extended DHCPv6 local server statistics on the specified logical system. If you do not specify a logical system, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Display information about extended DHCPv6 local server statistics on the specified routing instance. If you do not specify a routing instance, statistics are displayed for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">clear dhcpv6 server statistics on page 107</a></li></ul>
<b>List of Sample Output</b>	<a href="#">show dhcpv6 server statistics on page 133</a>
<b>Output Fields</b>	<a href="#">Table 8 on page 133</a> lists the output fields for the <b>show dhcpv6 server statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 8: show dhcpv6 server statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCPv6 local server because of errors. Only nonzero statistics appear in the Packets dropped output. When all of the Packets dropped statistics are 0 (zero), only the Total field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCPv6 local server</li> <li>• <b>Strict Reconfigure</b>—Number of solicit messages discarded because the client does not support reconfiguration</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error</li> <li>• <b>Send error</b>—Number of packets that the extended DHCPv6 local server could not send</li> </ul>
<b>Messages received</b>	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_CONFIRM</b>—Number of DHCPv6 CONFIRM PDUs received.</li> <li>• <b>DHCPV6_DECLINE</b>—Number of DHCPv6 DECLINE PDUs received.</li> <li>• <b>DHCPV6_INFORMATION_REQUEST</b>—Number of DHCPv6 INFORMATION-REQUEST PDUs received.</li> <li>• <b>DHCPV6_REBIND</b>—Number of DHCPv6 REBIND PDUs received.</li> <li>• <b>DHCPV6_RELAY_FORW</b>—Number of DHCPv6 RELAY-FORW PDUs received.</li> <li>• <b>DHCPV6_RELAY_REPL</b>—Number of DHCPv6 RELAY-REPL PDUs received.</li> <li>• <b>DHCPV6_RELEASE</b>—Number of DHCPv6 RELEASE PDUs received.</li> <li>• <b>DHCPV6_RENEW</b>—Number of DHCPv6 RENEW PDUs received.</li> <li>• <b>DHCPV6_REQUEST</b>—Number of DHCPv6 REQUEST PDUs received.</li> <li>• <b>DHCPV6_SOLICIT</b>—Number of DHCPv6 SOLICIT PDUs received.</li> </ul>
<b>Messages sent</b>	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> <li>• <b>DHCPV6_ADVERTISE</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHCPV6_REPLY</b>—Number of DHCPv6 ADVERTISE PDUs transmitted.</li> <li>• <b>DHC6_RECONFIGURE</b>—Number of DHCPv6 RECONFIGURE PDUs transmitted.</li> </ul>

## Sample Output

### show dhcpv6 server statistics

```
user@host> show dhcpv6 server statistics
```

Dhcpv6 Packets dropped:

Total	1
Lease Time Violation	1

Messages received:

DHCPV6_DECLINE	0
DHCPV6_SOLICIT	9
DHCPV6_INFORMATION_REQUEST	0
DHCPV6_RELEASE	0
DHCPV6_REQUEST	5
DHCPV6_CONFIRM	0
DHCPV6_RENEW	0
DHCPV6_REBIND	0
DHCPV6_RELAY_FORW	0
DHCPV6_RELAY_REPL	0

Messages sent:

DHCPV6_ADVERTISE	9
DHCPV6_REPLY	5
DHCPV6_RECONFIGURE	0

## CHAPTER 8

# Operational Commands (DHCP Relay Agent)

- clear dhcp relay binding
- clear dhcp relay statistics
- clear dhcpv6 relay binding
- clear dhcpv6 relay statistics
- show dhcp relay binding
- show dhcp relay statistics
- show dhcpv6 relay binding
- show dhcpv6 relay statistics
- show route extensive
- show route protocol

## clear dhcp relay binding

<b>Syntax</b>	<pre>clear dhcp relay binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Options <b>all</b> and <b>interface</b> added in Junos OS Release 8.4.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
<b>Description</b>	Clear the binding state of a Dynamic Host Configuration Protocol (DHCP) client from the client table.
<b>Options</b>	<p><b>address</b>—(Optional) Clear the binding state for the DHCP client, using one of the following entries:</p> <ul style="list-style-type: none"> <li><i>ip-address</i>—The specified IP address.</li> <li><i>mac-address</i>—The specified MAC address.</li> <li><i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCP clients.</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCP clients on the specified interface.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCP clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCP clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><i>Clearing DHCP Bindings for Subscriber Access</i></li> <li><a href="#">show dhcp relay binding on page 146</a></li> </ul>



**List of Sample Output** [clear dhcp relay binding on page 137](#)  
[clear dhcp relay binding all on page 137](#)  
[clear dhcp relay binding interface on page 137](#)  
[clear dhcp relay binding <interfaces-vlan> on page 137](#)  
[clear dhcp relay binding <interfaces-wildcard> on page 137](#)

**Output Fields** See [show dhcp relay binding](#) for an explanation of output fields.

## Sample Output

### clear dhcp relay binding

The following sample output displays the address bindings in the DHCP client table before and after the **clear dhcp relay binding** command is issued.

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
100.20.32.1     90:00:00:01:00:01 active    2007-02-08 16:41:17 EST
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

```
user@host> clear dhcp relay binding 100.20.32.1
```

```
user@host> show dhcp relay binding
IP address      Hardware address  Type    Lease expires at
192.168.14.8    90:00:01:01:02:01 active    2007-02-10 10:01:06 EST
```

### clear dhcp relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcp relay binding all
```

### clear dhcp relay binding interface

The following command clears DHCP relay agent bindings on a specific interface:

```
user@host> clear dhcp relay binding interface fe-0/0/3
```

### clear dhcp relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCP relay agent bindings on top of the underlying interface **ae0**, which clears DHCP bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcp relay binding interface ae0
```

### clear dhcp relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCP relay agent bindings over a specific interface:

```
user@host> clear dhcp relay binding ge-1/0/0.*
```

## clear dhcp relay statistics

---

Syntax	<code>clear dhcp relay statistics</code> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
Syntax	Syntax for EX Series switches:  <code>show dhcp relay statistics</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
Release Information	Command introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.1 for EX Series switches. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
Description	Clear all Dynamic Host Configuration Protocol (DHCP) relay statistics.
Options	<code>logical-system <i>logical-system-name</i></code> —(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.  <code>routing-instance <i>routing-instance-name</i></code> —(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"><li>• <a href="#">show dhcp relay statistics on page 151</a></li></ul>
List of Sample Output	<a href="#">clear dhcp relay statistics on page 139</a>
Output Fields	<a href="#">Table 9 on page 139</a> lists the output fields for the <code>clear dhcp relay statistics</code> command.

Table 9: clear dhcp relay statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP relay agent application.</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified.</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified.</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified.</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment.</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface.</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance.</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address.</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short.</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error.</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP relay application could not send.</li> <li>• <b>Option 60</b>—Number of packets discarded containing DHCP option 60 vendor-specific information.</li> <li>• <b>Option 82</b>—Number of packets discarded because DHCP option 82 information could not be added.</li> </ul>
<b>Messages received</b>	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHC PNACK</b>—Number of DHCP NACK PDUs transmitted</li> </ul>

## Sample Output

### clear dhcp relay statistics

The following sample output displays the DHCP relay statistics before and after the **clear dhcp relay statistics** command is issued.

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
  Total          1
  Lease Time Violated 1

Messages received:
  BOOTREQUEST    116
  DHCPDECLINE    0
  DHCPDISCOVER   11
  DHCPINFORM     0
  DHCPRELEASE    0
  DHCPREQUEST    105

Messages sent:
  BOOTREPLY      44
  DHCPOFFER      11
  DHCPACK        11
  DHCPNAK        11
```

```
user@host> clear dhcp relay statistics
```

```
user@host> show dhcp relay statistics
```

```
Packets dropped:
  Total          0

Messages received:
  BOOTREQUEST    0
  DHCPDECLINE    0
  DHCPDISCOVER   0
  DHCPINFORM     0
  DHCPRELEASE    0
  DHCPREQUEST    0

Messages sent:
  BOOTREPLY      0
  DHCPOFFER      0
  DHCPACK        0
  DHCPNAK        0
```

## clear dhcpv6 relay binding

<b>Syntax</b>	<pre>clear dhcpv6 relay binding &lt;address&gt; &lt;all&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 11.4.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Options <i>interfaces-vlan</i> and <i>interfaces-wildcard</i> added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
<b>Description</b>	Clear the binding state of Dynamic Host Configuration Protocol for IPv6 (DHCPv6) clients from the client table.
<b>Options</b>	<p><b>address</b>—(Optional) Clear the binding state for the DHCPv6 client, using one of the following entries:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>all</b>—(Optional) Clear the binding state for all DHCPv6 clients.</p> <p><b>interfaces-vlan</b>—(Optional) Clear the binding state on the interface VLAN ID and S-VLAN ID.</p> <p><b>interfaces-wildcard</b>—(Optional) The set of interfaces on which to clear bindings. This option supports the use of the wildcard character (*).</p> <p><b>interface interface-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified interface.</p> <p><b>logical-system logical-system-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Clear the binding state for DHCPv6 clients on the specified routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Clearing DHCP Bindings for Subscriber Access</i></li> <li>• <a href="#">show dhcpv6 relay binding on page 154</a></li> </ul>

- List of Sample Output**
- [clear dhcpv6 relay binding on page 142](#)
  - [clear dhcpv6 relay binding <prefix> on page 142](#)
  - [clear dhcpv6 relay binding all on page 142](#)
  - [clear dhcv6p relay binding interface on page 142](#)
  - [clear dhcpv6 relay binding <interfaces-vlan> on page 143](#)
  - [clear dhcpv6 relay binding <interfaces-wildcard> on page 143](#)

**Output Fields** See [show dhcpv6 relay binding](#) for an explanation of output fields.

## Sample Output

### clear dhcpv6 relay binding

The following sample output displays the DHCPv6 bindings before and after the **clear dhcpv6 relay binding** command is issued.

```
user@host> show dhcpv6 relay binding
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:bd8:3c4d:15::/64	1	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01					
2001:bd8:3c4d:16::/64	2	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02					
2001:bd8:3c4d:17::/64	3	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03					
2001:bd8:3c4d:18::/64	4	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04					
2001:bd8:3c4d:19::/64	5	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05					
2001:bd8:3c4d:20::/64	6	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06					

### clear dhcpv6 relay binding <prefix>

```
user@host> clear dhcpv6 relay binding 2001:bd8:3c4d:15::/64
```

```
user@host> show dhcpv6 relay binding
```

Prefix	Session Id	Expires	State	Interface	Client DUID
2001:bd8:3c4d:16::/64	2	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02					
2001:bd8:3c4d:17::/64	3	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03					
2001:bd8:3c4d:18::/64	4	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04					
2001:bd8:3c4d:19::/64	5	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05					
2001:bd8:3c4d:20::/64	6	83720	BOUND	ge-1/0/0.0	
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06					

### clear dhcpv6 relay binding all

The following command clears all DHCP relay agent bindings:

```
user@host> clear dhcpv6 relay binding all
```

### clear dhcv6p relay binding interface

The following command clears DHCPv6 relay agent bindings on a specific interface:

```
user@host> clear dhcpv6 relay binding interface fe-0/0/2
```

#### clear dhcpv6 relay binding <interfaces-vlan>

The following command uses the *interfaces-vlan* option to clear all DHCPv6 relay agent bindings on top of the underlying interface **ae0**, which clears DHCPv6 bindings on all demux VLANs on top of **ae0**:

```
user@host> clear dhcpv6 relay binding interface ae0
```

#### clear dhcpv6 relay binding <interfaces-wildcard>

The following command uses the *interfaces-wildcard* option to clear all DHCPv6 relay agent bindings over a specific interface:

```
user@host> clear dhcpv6 relay binding ge-1/0/0.*
```

## clear dhcpv6 relay statistics

<b>Syntax</b>	<b>clear dhcpv6 relay statistics</b> <b>&lt;logical-system <i>logical-system-name</i>&gt;</b> <b>&lt;routing-instance <i>routing-instance-name</i>&gt;</b>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.
<b>Description</b>	Clear all Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are cleared for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are cleared for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">clear dhcpv6 relay statistics on page 144</a>
<b>Output Fields</b>	See <a href="#">show dhcpv6 relay statistics</a> for an explanation of output fields.

## Sample Output

### clear dhcpv6 relay statistics

The following sample output displays the DHCPv6 relay statistics before and after the **clear dhcpv6 relay statistics** command is issued.

```

user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total                0
    Lease Time Violated  1

Messages received:
    DHCPV6_DECLINE        0
    DHCPV6_SOLICIT        10
    DHCPV6_INFORMATION_REQUEST  0
    DHCPV6_RELEASE        0
    DHCPV6_REQUEST        10
    DHCPV6_CONFIRM        0
    DHCPV6_RENEW          0
    DHCPV6_REBIND         0
    DHCPV6_RELAY_REPL     0

Messages sent:
    DHCPV6_ADVERTISE      0
    DHCPV6_REPLY           0
    DHCPV6_RECONFIGURE    0
    DHCPV6_RELAY_FORW     0

```



```
user@host> clear dhcpv6 relay statistics
```

```
user@host> show dhcpv6 relay statistics
```

```
DHCPv6 Packets dropped:
```

```
    Total                                0
```

```
Messages received:
```

```
    DHCPV6_DECLINE                      0
```

```
    DHCPV6_SOLICIT                      0
```

```
    DHCPV6_INFORMATION_REQUEST          0
```

```
    DHCPV6_RELEASE                      0
```

```
    DHCPV6_REQUEST                      0
```

```
    DHCPV6_CONFIRM                      0
```

```
    DHCPV6_RENEW                        0
```

```
    DHCPV6_REBIND                       0
```

```
    DHCPV6_RELAY_REPL                   0
```

```
Messages sent:
```

```
    DHCPV6_ADVERTISE                    0
```

```
    DHCPV6_REPLY                        0
```

```
    DHCPV6_RECONFIGURE                  0
```

```
    DHCPV6_RELAY_FORW                   0
```

## show dhcp relay binding

---

**Syntax**    **show dhcp relay binding**  
              <address>  
              <brief>  
              <detail>  
              <interface *interface-name*>  
              <interfaces-vlan>  
              <interfaces-wildcard>  
              <ip-address | mac-address>  
              <logical-system *logical-system-name*>  
              <routing-instance *routing-instance-name*>  
              <summary>

**Release Information**    Command introduced in Junos OS Release 8.3.  
                              Options **interface** and **mac-address** added in Junos OS Release 8.4.  
                              Options **interfaces-vlan** and **interfaces-wildcard** added in Junos OS Release 12.1.  
                              Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.

**Description**    Display the address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.

**Options**    **address**—(Optional) Display DHCP binding information for a specific client identified by one of the following entries:

- *ip-address*—The specified IP address.
- *mac-address*—The specified MAC address.
- *session-id*—The specified session ID.

**brief**—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as **show dhcp relay binding**.

**detail**—(Optional) Display detailed client binding information.

**interface *interface-name***—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and SVLAN ID.

**interfaces-vlan**—(Optional) Show the binding state information on the interface VLAN ID and S-VLAN ID.

**interfaces-wildcard**—(Optional) The set of interfaces on which to show binding state information. This option supports the use of the wildcard character (\*).

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

**routing-instance *routing-instance-name***—(Optional) Perform this operation on the specified routing instance.

**summary**—(Optional) Display a summary of DHCP client information.

**Required Privilege Level** view

**Related Documentation**

- [Clearing DHCP Bindings for Subscriber Access](#)
- [clear dhcp relay binding on page 136](#)

**List of Sample Output**

- [show dhcp relay binding on page 148](#)
- [show dhcp relay binding detail on page 149](#)
- [show dhcp relay binding interface on page 149](#)
- [show dhcp relay binding interface vlan-id on page 149](#)
- [show dhcp relay binding interface svlan-id on page 149](#)
- [show dhcp relay binding ip-address on page 150](#)
- [show dhcp relay binding mac-address on page 150](#)
- [show dhcp relay binding session-id on page 150](#)
- [show dhcp relay binding <interfaces-vlan> on page 150](#)
- [show dhcp relay binding <interfaces-wildcard> on page 150](#)
- [show dhcp relay binding summary on page 150](#)

**Output Fields** Table 10 on page 147 lists the output fields for the **show dhcp relay binding** command. Output fields are listed in the approximate order in which they appear.

**Table 10: show dhcp relay binding Output Fields**

Field Name	Field Description	Level of Output
<i>number</i> clients, ( <i>number</i> init, <i>number</i> bound, <i>number</i> selecting, <i>number</i> requesting, <i>number</i> renewing, <i>number</i> rebinding, <i>number</i> releasing)	Summary counts of the total number of DHCP clients and the number of DHCP clients in each state.	<b>summary</b>
IP address	IP address of the DHCP client.	<b>briefdetail</b>
Session Id	Session ID of the subscriber session.	<b>briefdetail</b>
Generated Remote ID	Remote ID generated by the Option 82 Agent Remote ID (suboption 1)	<b>detail</b>
Hardware address	Hardware address of the DHCP client.	<b>briefdetail</b>
Expires	Number of seconds in which the lease expires.	<b>briefdetail</b>

Table 10: show dhcp relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the DHCP relay address binding table on the DHCP client: <ul style="list-style-type: none"> <li><b>BOUND</b>—Client has an active IP address lease.</li> <li><b>INIT</b>—Initial state.</li> <li><b>REBINDING</b>—Client is broadcasting a request to renew the IP address lease.</li> <li><b>RELEASE</b>—Client is releasing the IP address lease.</li> <li><b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li><b>REQUESTING</b>—Client is requesting a DHCP server.</li> <li><b>SELECTING</b>—Client is receiving offers from DHCP servers.</li> </ul>	<b>brief</b> <b>detail</b>
<b>Interface</b>	Incoming client interface.	<b>brief</b>
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>
<b>Lease Expires in</b>	Number of seconds in which the lease expires.	<b>detail</b>
<b>Lease Start</b>	Date and time at which the client's IP address lease started.	<b>detail</b>
<b>Lease time violated</b>	Lease time violation has occurred.	<b>detail</b>
<b>Incoming Client Interface</b>	Client's incoming interface.	<b>detail</b>
<b>Server IP Address</b>	IP address of the DHCP server.	<b>detail</b>
<b>Server Interface</b>	Interface of the DHCP server.	<b>detail</b>
<b>Bootp Relay Address</b>	IP address of BOOTP relay.	<b>detail</b>
<b>Type</b>	Type of DHCP packet processing performed on the router: <ul style="list-style-type: none"> <li><b>active</b>—Router actively processes and relays DHCP packets.</li> <li><b>passive</b>—Router passively snoops DHCP packets passing through the router.</li> </ul>	<b>All levels</b>
<b>Lease expires at</b>	Date and time at which the client's IP address lease expires.	<b>All levels</b>

## Sample Output

### show dhcp relay binding

```

user@host> show dhcp relay binding
IP address      Session Id  Hardware address  Expires   State   Interface
100.20.32.11    41         00:10:94:00:00:01 86371     BOUND   ge-1/0/0.0
100.20.32.12    42         00:10:94:00:00:02 86371     BOUND   ge-1/0/0.0

```

100.20.32.13	43	00:10:94:00:00:03	86371	BOUND	ge-1/0/0.0
100.20.32.14	44	00:10:94:00:00:04	86371	BOUND	ge-1/0/0.0
100.20.32.15	45	00:10:94:00:00:05	86371	BOUND	ge-1/0/0.0

### show dhcp relay binding detail

```
user@host> show dhcp relay binding detail
```

```
Client IP Address: 100.20.32.11
  Hardware Address: 00:10:94:00:00:01
  State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
  Lease Expires: 2009-07-21 11:00:06 PDT
  Lease Expires in: 86361 seconds
  Lease Start: 2009-07-20 11:00:06 PDT
  Lease time violated: yes
  Last Packet Received: 2009-07-20 11:00:06 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address: 100.20.22.2
  Server Interface: none
  Bootp Relay Address: 100.20.32.2
  Session Id: 41
```

```
Client IP Address: 100.20.32.12
  Hardware Address: 00:10:94:00:00:02
  State: BOUND(DHCP_RELAY_STATE_BOUND_ON_INTF_DELETE)
  Lease Expires: 2009-07-21 11:00:06 PDT
  Lease Expires in: 86361 seconds
  Lease Start: 2009-07-20 11:00:06 PDT
  Last Packet Received: 2009-07-20 11:00:06 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Ip Address: 100.20.22.2
  Server Interface: none
  Bootp Relay Address: 100.20.32.2
  Session Id: 42
  Generated Remote ID: host:ge-1/0/0:100
```

### show dhcp relay binding interface

```
user@host> show dhcp relay binding interface fe-0/0/2
```

IP address	Hardware address	Type	Lease expires at
100.20.32.1	90:00:00:01:00:01	active	2007-03-27 15:06:20 EDT

### show dhcp relay binding interface vlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:100
```

IP address	Session Id	Hardware address	Expires	State	Interface
200.20.20.15	6	00:10:94:00:00:01	86124	BOUND	ge-1/1/0:100

### show dhcp relay binding interface svlan-id

```
user@host> show dhcp relay binding interface ge-1/1/0:10-100
```

IP address	Session Id	Hardware address	Expires	State	Interface
------------	------------	------------------	---------	-------	-----------

```

200.20.20.16      7          00:10:94:00:00:02  86124      BOUND
ge-1/1/0:10-100

```

#### show dhcp relay binding ip-address

```

user@host> show dhcp relay binding 100.20.32.13
IP address      Session Id  Hardware address  Expires    State      Interface
100.20.32.13    43         00:10:94:00:00:03  86293     BOUND      ge-1/0/0.0

```

#### show dhcp relay binding mac-address

```

user@host> show dhcp relay binding 00:10:94:00:00:05
IP address      Session Id  Hardware address  Expires    State      Interface
100.20.32.15    45         00:10:94:00:00:05  86279     BOUND      ge-1/0/0.0

```

#### show dhcp relay binding session-id

```

user@host> show dhcp relay binding 41
IP address      Session Id  Hardware address  Expires    State      Interface
100.20.32.11    41         00:10:94:00:00:01  86305     BOUND      ge-1/0/0.0

```

#### show dhcp relay binding <interfaces-vlan>

```

user@host> show dhcp relay binding ge-1/0/0:100-200
IP address      Session Id  Hardware address  Expires    State      Interface
192.168.0.17    42         00:10:94:00:00:02  86346     BOUND      ge-1/0/0.1073741827
192.168.0.16    41         00:10:94:00:00:01  86346     BOUND      ge-1/0/0.1073741827

```

#### show dhcp relay binding <interfaces-wildcard>

```

user@host> show dhcp relay binding ge-1/3/*
IP address      Session Id  Hardware address  Expires    State      Interface
192.168.0.9     24         00:10:94:00:00:04  86361     BOUND      ge-1/3/0.110
192.168.0.8     23         00:10:94:00:00:03  86361     BOUND      ge-1/3/0.110
192.168.0.7     22         00:10:94:00:00:02  86361     BOUND      ge-1/3/0.110

```

#### show dhcp relay binding summary

```

user@host> show dhcp relay binding summary
3 clients, (2 init, 1 bound, 0 selecting, 0 requesting, 0 renewing, 0 rebinding,
0 releasing)

```

## show dhcp relay statistics

---

<b>Syntax</b>	<pre>show dhcp relay statistics &lt;logical-system <i>logical-system-name</i>&gt; &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Syntax</b>	<p>Syntax for EX Series switches:</p> <pre>show dhcp relay statistics &lt;routing-instance <i>routing-instance-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Routers.</p>
<b>Description</b>	Display Dynamic Host Configuration Protocol (DHCP) relay statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(On routers only) (Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear dhcp relay statistics on page 138</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show dhcp relay statistics on page 153</a>
<b>Output Fields</b>	<p><a href="#">Table 11 on page 152</a> lists the output fields for the <b>show dhcp relay statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 11: show dhcp relay statistics Output Fields

Field Name	Field Description
<b>Packets dropped</b>	<p>Number of packets discarded by the extended DHCP relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li>• <b>Total</b>—Total number of packets discarded by the extended DHCP relay agent application.</li> <li>• <b>Bad hardware address</b>—Number of packets discarded because an invalid hardware address was specified.</li> <li>• <b>Bad opcode</b>—Number of packets discarded because an invalid operation code was specified.</li> <li>• <b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li>• <b>Invalid server address</b>—Number of packets discarded because an invalid server address was specified.</li> <li>• <b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li>• <b>No available addresses</b>—Number of packets discarded because there were no addresses available for assignment.</li> <li>• <b>No interface match</b>—Number of packets discarded because they did not belong to a configured interface.</li> <li>• <b>No routing instance match</b>—Number of packets discarded because they did not belong to a configured routing instance.</li> <li>• <b>No valid local address</b>—Number of packets discarded because there was no valid local address.</li> <li>• <b>Packet too short</b>—Number of packets discarded because they were too short.</li> <li>• <b>Read error</b>—Number of packets discarded because of a system read error.</li> <li>• <b>Send error</b>—Number of packets that the extended DHCP relay application could not send.</li> <li>• <b>Option 60</b>—Number of packets discarded containing DHCP option 60 vendor-specific information.</li> <li>• <b>Option 82</b>—Number of packets discarded because DHCP option 82 information could not be added.</li> </ul>
<b>Messages received</b>	<p>Number of DHCP messages received.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTP protocol data units (PDUs) received</li> <li>• <b>DHCPDECLINE</b>—Number of DHCP PDUs of type DECLINE received</li> <li>• <b>DHCPDISCOVER</b>—Number of DHCP PDUs of type DISCOVER received</li> <li>• <b>DHCPINFORM</b>—Number of DHCP PDUs of type INFORM received</li> <li>• <b>DHCPRELEASE</b>—Number of DHCP PDUs of type RELEASE received</li> <li>• <b>DHCPREQUEST</b>—Number of DHCP PDUs of type REQUEST received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCP messages sent.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREPLY</b>—Number of BOOTP PDUs transmitted</li> <li>• <b>DHCPOFFER</b>—Number of DHCP OFFER PDUs transmitted</li> <li>• <b>DHCPACK</b>—Number of DHCP ACK PDUs transmitted</li> <li>• <b>DHCPNACK</b>—Number of DHCP NACK PDUs transmitted</li> <li>• <b>DHCPFORCERENEW</b>—Number of DHCP FORCERENEW PDUs transmitted</li> </ul>
<b>External Server Response</b>	<p>State of the external DHCP server responsiveness.</p>
<b>Packets forwarded</b>	<p>Number of packets forwarded.</p> <ul style="list-style-type: none"> <li>• <b>BOOTREQUEST</b>—Number of BOOTREQUEST protocol data units (PDUs) forwarded</li> <li>• <b>BOOTREPLY</b>—Number of BOOTREPLY protocol data units (PDUs) forwarded</li> </ul>



Table 11: show dhcp relay statistics Output Fields (*continued*)

Field Name	Field Description
External Server Response	State of the external DHCP server responsiveness.

## Sample Output

### show dhcp relay statistics

```

user@host> show dhcp relay statistics
Packets dropped:
    Total                34
    Bad hardware address  1
    Bad opcode            1
    Bad options           3
    Invalid server address 5
    Lease Time Violation  1
    No available addresses 1
    No interface match    2
    No routing instance match 9
    No valid local address 4
    Packet too short      2
    Read error            1
    Send error            1
    Option 60             1
    Option 82             2

Messages received:
    BOOTREQUEST          116
    DHCPDECLINE           0
    DHCPDISCOVER          11
    DHCPINFORM            0
    DHCPRELEASE           0
    DHCPREQUEST          105

Messages sent:
    BOOTREPLY            0
    DHCPOFFER             2
    DHCPACK               1
    DHCPNAK               0
    DHCPFORCERENEW        0

Packets forwarded:
    Total                4
    BOOTREQUEST           2
    BOOTREPLY             2

External Server Response:
    State                Responding

```

## show dhcpv6 relay binding

<b>Syntax</b>	<pre>show dhcpv6 relay binding &lt;address&gt; &lt;brief&gt; &lt;detail&gt; &lt;interface interface-name&gt; &lt;interfaces-vlan&gt; &lt;interfaces-wildcard&gt; &lt;logical-system logical-system-name&gt; &lt;routing-instance routing-instance-name&gt; &lt;summary&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 11.4.</p> <p><i>interfaces-vlan</i> and <i>interfaces-wildcard</i> options introduced in Junos OS Release 12.1.</p>
<b>Description</b>	Display the DHCPv6 address bindings in the Dynamic Host Configuration Protocol (DHCP) client table.
<b>Options</b>	<p><b>address</b>—(Optional) One of the following identifiers for the DHCPv6 client whose binding state you want to show:</p> <ul style="list-style-type: none"> <li>• <i>CID</i>—The specified Client ID (CID).</li> <li>• <i>ipv6-prefix</i>—The specified IPv6 prefix.</li> <li>• <i>session-id</i>—The specified session ID.</li> </ul> <p><b>brief</b>—(Optional) Display brief information about the active client bindings. This is the default, and produces the same output as <b>show dhcpv6 relay binding</b>.</p> <p><b>detail</b>—(Optional) Display detailed client binding information.</p> <p><b>interface interface-name</b>—(Optional) Perform this operation on the specified interface. You can optionally filter on VLAN ID and S-VLAN ID.</p> <p><b>interfaces-vlan</b>—(Optional) Interface VLAN ID or S-VLAN ID interface on which to show binding state information.</p> <p><b>interfaces-wildcard</b>—(Optional) Set of interfaces on which to show binding state information. This option supports the use of the wildcard character (*).</p> <p><b>logical-system logical-system-name</b>—(Optional) Perform this operation on the specified logical system.</p> <p><b>routing-instance routing-instance-name</b>—(Optional) Perform this operation on the specified routing instance.</p> <p><b>summary</b>—(Optional) Display a summary of DHCPv6 client information.</p>
<b>Required Privilege Level</b>	view

- Related Documentation**
- [Clearing DHCP Bindings for Subscriber Access](#)
  - [clear dhcpv6 relay binding on page 141](#)

- List of Sample Output**
- [show dhcpv6 relay binding on page 156](#)
  - [show dhcpv6 relay binding \(Address\) on page 157](#)
  - [show dhcpv6 relay binding detail \(Client ID\) on page 157](#)
  - [show dhcpv6 relay binding detail on page 157](#)
  - [show dhcpv6 relay binding detail \(Multi-Relay Topology\) on page 158](#)
  - [show dhcpv6 relay binding \(Session ID\) on page 158](#)
  - [show dhcpv6 relay binding \(Interfaces VLAN\) on page 158](#)
  - [show dhcpv6 relay binding \(Interfaces Wildcard\) on page 158](#)
  - [show dhcpv6 relay binding \(Interfaces Wildcard\) on page 159](#)
  - [show dhcpv6 relay binding summary on page 159](#)

**Output Fields** Table 12 on page 155 lists the output fields for the **show dhcpv6 relay binding** command. Output fields are listed in the approximate order in which they appear.

**Table 12: show dhcpv6 relay binding Output Fields**

Field Name	Field Description	Level of Output
<i>number clients, (number init, number bound, number selecting, number requesting, number renewing, number rebinding, number releasing)</i>	Summary counts of the total number of DHCPv6 clients and the number of DHCPv6 clients in each state.	<b>summary</b>
<b>Client IPv6 Prefix</b>	Prefix of the DHCPv6 client.	<b>brief detail</b>
<b>Client DUID</b>	DHCP for IPv6 Unique Identifier (DUID) of the client.	<b>brief detail</b>
<b>Session Id</b>	Session ID of the subscriber session.	<b>brief detail</b>
<b>Expires</b>	Number of seconds in which the lease expires.	<b>brief detail</b>
<b>State</b>	State of the DHCPv6 relay address binding table on the DHCPv6 client: <ul style="list-style-type: none"> <li>• <b>BOUND</b>—Client has an active IP address lease.</li> <li>• <b>INIT</b>—Initial state.</li> <li>• <b>REBINDING</b>—Client is broadcasting a request to renew the IP address lease.</li> <li>• <b>RELEASE</b>—Client is releasing the IP address lease.</li> <li>• <b>RENEWING</b>—Client is sending a request to renew the IP address lease.</li> <li>• <b>REQUESTING</b>—Client is requesting a DHCPv6 server.</li> <li>• <b>SELECTING</b>—Client is receiving offers from DHCPv6 servers.</li> </ul>	<b>brief detail</b>
<b>Interface</b>	Incoming client interface.	<b>brief</b>
<b>Lease Expires</b>	Date and time at which the client's IP address lease expires.	<b>detail</b>

Table 12: show dhcpv6 relay binding Output Fields (*continued*)

Field Name	Field Description	Level of Output
Lease Expires in	Number of seconds in which the lease expires.	detail
Preferred Lease Expires	Date and UTC time at which the client's IPv6 prefix expires.	detail
Preferred Lease Expires in	Number of seconds at which the client's IPv6 prefix expires.	detail
Lease Start	Date and time at which the client's IP address lease started.	detail
Lease time violated	Lease time violation has occurred.	detail
Incoming Client Interface	Client's incoming interface.	detail
Server Address	IP address of the DHCPv6 server.  Displays <b>unknown</b> for a DHCPv6 relay agent in a multi-relay topology that is not directly adjacent to the DHCPv6 server and does not detect the IP address of the server. In that case, the output instead displays the <b>Next Hop Server Facing Relay</b> field.	detail
Next Hop Server Facing Relay	Next-hop address in the direction of the DHCPv6 server.	detail
Server Interface	Interface of the DHCPv6 server.	detail
Relay Address	IP address of the relay.	detail
Client Pool Name	Address pool that granted the client lease.	detail
Client ID Length	Length of client ID.	All levels
Client Id	Client ID.	All levels
Generated Circuit ID	Circuit ID generated by the DHCPv6 Interface-ID option (option 18)	detail
Generated Remote ID Enterprise Number	The Juniper Networks IANA private enterprise number	detail
Generated Remote ID	Remote ID generated by the DHCPv6 Remote-ID option (option 37)	detail

## Sample Output

### show dhcpv6 relay binding

```

user@host> show dhcpv6 relay binding
Prefix          Session Id  Expires  State  Interface  Client DUID
2001:bd8:3c4d:15::/64  1          83720    BOUND  ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01
2001:bd8:3c4d:16::/64  2          83720    BOUND  ge-1/0/0.0

```

```

LL_TIME0x1-0x4bfa26af-00:10:94:00:00:02
2001:bd8:3c4d:17::/64      3      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:03
2001:bd8:3c4d:18::/64      4      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:04
2001:bd8:3c4d:19::/64      5      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:05
2001:bd8:3c4d:20::/64      6      83720      BOUND      ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:06

```

### show dhcpv6 relay binding (Address)

```

user@host> show dhcpv6 relay binding 2001:bd8:1111:2222::/64 detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
  Client DUID:              LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                    BOUND(RELAY_STATE_BOUND)
  Lease Expires:            2011-05-25 07:12:09 PDT
  Lease Expires in:         77115 seconds
  Preferred Lease Expires:  2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:              2011-05-24 07:12:09 PDT
  Incoming Client Interface: ge-1/0/0.0
  Server Address:           2008:aaaa:bbbb::1
  Server Interface:         none
  Relay Address:            2001:bd8:1111:2222::
  Client Pool Name:         pool-25
  Client Id Length:         14
  Client Id:
    /0x00010001/0x4bfa26af/0x00109400/0x0001

```

### show dhcpv6 relay binding detail (Client ID)

```

user@host> show dhcpv6 relay binding 14/0x00010001/0x4bfa26af/0x00109400/0x0001
detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64
  Client DUID:              LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

  State:                    BOUND(RELAY_STATE_BOUND)
  Lease Expires:            2011-05-25 07:12:09 PDT
  Lease Expires in:         77115 seconds
  Preferred Lease Expires:  2012-07-24 00:18:14 UTC
  Preferred Lease Expires in: 600 seconds
  Lease Start:              2011-05-24 07:12:09 PDT
  Lease time violated:      yes
  Incoming Client Interface: ge-1/0/0.0
  Server Address:           2008:aaaa:bbbb::1
  Server Interface:         none
  Relay Address:            2001:bd8:1111:2222::
  Client Pool Name:         pool-25
  Client Id Length:         14
  Client Id:
    /0x00010001/0x4bfa26af/0x00109400/0x0001

```

### show dhcpv6 relay binding detail

```

user@host> show dhcpv6 relay binding detail
Session Id: 1
  Client IPv6 Prefix:      2001:bd8:3c4d:15::/64

```

```

Client DUID:                               LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

State:                                     BOUND(RELAY_STATE_BOUND)
Lease Expires:                             2011-05-25 07:12:09 PDT
Lease Expires in:                           77115 seconds
Preferred Lease Expires:                     2012-07-24 00:18:14 UTC
Preferred Lease Expires in:                   600 seconds
Lease Start:                                2011-05-24 07:12:09 PDT
Lease time violated:                          yes
Incoming Client Interface:                   ge-1/0/0.0
Server Address:                              2008:aaaa:bbbb::1
Server Interface:                            none
Relay Address:                               2001:bd8:1111:2222::
Client Pool Name:                            pool-25
Client Id Length:                            14
Client Id:                                   /0x00010001/0x4bfa26af/0x00109400/0x0001
Generated Remote ID Enterprise Number:       1411
Generated Remote ID:                         host:ge-1/0/0:100

```

#### show dhcpv6 relay binding detail (Multi-Relay Topology)

```

user@host > show dhcpv6 relay binding detail
Session Id: 13
Client IPv6 Prefix:                        3000:0:0:8001::5/128
Client DUID:                               LL0x1-00:00:65:03:01:02
State:                                     BOUND(DHCPV6_RELAY_STATE_BOUND)
Lease Expires:                             2011-11-21 06:14:50 PST
Lease Expires in:                           293 seconds
Preferred Lease Expires:                     2012-07-24 00:18:14 UTC
Preferred Lease Expires in:                   600 seconds
Lease Start:                                2011-11-21 06:09:50 PST
Incoming Client Interface:                   ge-1/0/0.0
Server Address:                              unknown
Next Hop Server Facing Relay:                4000::2
Server Interface:                            none
Client Id Length:                            10
Client Id:                                   /0x00030001/0x00006503/0x0102

```

#### show dhcpv6 relay binding (Session ID)

```

user@host> show dhcpv6 relay binding 41
Prefix          Session Id Expires   State   Interface   Client DUID
2001:bd8:3c4d:15::/64  41      78837   BOUND   ge-1/0/0.0
LL_TIME0x1-0x4bfa26af-00:10:94:00:00:01

```

#### show dhcpv6 relay binding (Interfaces VLAN)

```

user@host> show dhcpv6 relay binding ge-1/0/0:100-200
Prefix          Session Id Expires   State   Interface   Client DUID
2001:DB8::/32   11        87583   BOUND   ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   12        87583   BOUND   ge-1/0/0.1073741827
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

#### show dhcpv6 relay binding (Interfaces Wildcard)

```

user@host> show dhcpv6 relay binding demux0
Prefix          Session Id Expires   State   Interface   Client DUID
2001:DB8::/32   30        79681   BOUND   demux0.1073741824
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   31        79681   BOUND   demux0.1073741825

```

```

LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32      32      79681    BOUND    demux0.1073741826
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

#### show dhcpv6 relay binding (Interfaces Wildcard)

```

user@host> show dhcpv6 relay binding ge-1/3/*
Prefix          Session Id Expires State Interface Client DUID
2001:DB8::/32   22      79681    BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:DB9::/32   33      79681    BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01
2001:CB9::/32   24      79681    BOUND ge-1/3/0.110
LL_TIME0x1-0x4d5d009f-00:10:94:00:00:01

```

#### show dhcpv6 relay binding summary

```

user@host> show dhcpv6 relay binding summary
5 clients, (0 init, 5 bound, 0 selecting, 0 requesting, 0 renewing, 0 releasing)

```

## show dhcpv6 relay statistics

<b>Syntax</b>	<b>show dhcpv6 relay statistics</b> <code>&lt;logical-system <i>logical-system-name</i>&gt;</code> <code>&lt;routing-instance <i>routing-instance-name</i>&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 12.1X48R3 for PTX Series Packet Transport Switches. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
<b>Description</b>	Display Dynamic Host Configuration Protocol for IPv6 (DHCPv6) relay statistics.
<b>Options</b>	<p><b>logical-system <i>logical-system-name</i></b>—(Optional) Perform this operation on the specified logical system. If you do not specify a logical system name, statistics are displayed for the default logical system.</p> <p><b>routing-instance <i>routing-instance-name</i></b>—(Optional) Perform this operation on the specified routing instance. If you do not specify a routing instance name, statistics are displayed for the default routing instance.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear dhcpv6 relay statistics on page 144</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show dhcpv6 relay statistics on page 161</a>
<b>Output Fields</b>	Table 13 on page 160 lists the output fields for the <b>show dhcpv6 relay statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 13: show dhcpv6 relay statistics Output Fields

Field Name	Field Description
DHCPv6 Packets dropped	<p>Number of packets discarded by the extended DHCPv6 relay agent application due to errors. Only nonzero statistics appear in the <b>Packets dropped</b> output. When all of the Packets dropped statistics are 0 (zero), only the <b>Total</b> field appears.</p> <ul style="list-style-type: none"> <li><b>Total</b>—Total number of packets discarded by the DHCPv6 relay agent application.</li> <li><b>Bad options</b>—Number of packets discarded because invalid options were specified.</li> <li><b>Bad send</b>—Number of packets that the extended DHCP relay application could not send.</li> <li><b>Bad src address</b>—Number of packets discarded because the family type was not AF_INET6.</li> <li><b>No client id</b>—Number of packets discarded because they could not be matched to a client.</li> <li><b>Lease Time Violation</b>—Number of packets discarded because of a lease time violation</li> <li><b>No safd</b>—Number of packets discarded because they arrived on an unconfigured interface.</li> <li><b>Short packet</b>—Number of packets discarded because they were too short.</li> <li><b>Relay hop count</b>—Number of packets discarded because the hop count in the packet exceeded 32.</li> </ul>



Table 13: show dhcpv6 relay statistics Output Fields (*continued*)

Field Name	Field Description
<b>Messages received</b>	<p>Number of DHCPv6 messages received.</p> <ul style="list-style-type: none"> <li><b>DHCPV6_DECLINE</b>—Number of DHCPv6 PDUs of type DECLINE received</li> <li><b>DHCPV6_SOLICIT</b>—Number of DHCPv6 PDUs of type SOLICIT received</li> <li><b>DHCPV6_INFORMATION_REQUEST</b>—Number of DHCPv6 PDUs of type INFORMATION-REQUEST received</li> <li><b>DHCPV6_RELEASE</b>—Number of DHCPv6 PDUs of type RELEASE received</li> <li><b>DHCPV6_REQUEST</b>—Number of DHCPv6 PDUs of type REQUEST received</li> <li><b>DHCPV6_CONFIRM</b>—Number of DHCPv6 PDUs of type CONFIRM received</li> <li><b>DHCPV6_RENEW</b>—Number of DHCPv6 PDUs of type RENEW received</li> <li><b>DHCPV6_REBIND</b>—Number of DHCPv6 PDUs of type REBIND received</li> <li><b>DHCPV6_RELAY_REPL</b>—Number of DHCPv6 PDUs of type RELAY-REPL received</li> </ul>
<b>Messages sent</b>	<p>Number of DHCPv6 messages sent.</p> <ul style="list-style-type: none"> <li><b>DHCPV6_ADVERTISE</b>—Number of DHCPv6 ADVERTISE PDUs transmitted</li> <li><b>DHCP_REPLY</b>—Number of DHCPv6 REPLY PDUs transmitted</li> <li><b>DHCP_RECONFIGURE</b>—Number of DHCPv6 RECONFIGURE PDUs transmitted</li> <li><b>DHCP_RELAY_FORW</b>—Number of DHCPv6 RELAY-FORW PDUs transmitted</li> </ul>
<b>Packets forwarded</b>	<p>Number of packets forwarded by the extended DHCPv6 relay agent application.</p> <ul style="list-style-type: none"> <li><b>FWD REQUEST</b>—Number of DHCPv6 REQUEST packets forwarded</li> <li><b>FWD REPLY</b>—Number of DHCPv6 REPLY packets forwarded</li> </ul>
<b>External Server Response</b>	State of the external DHCP server responsiveness.

## Sample Output

### show dhcpv6 relay statistics

```

user@host> show dhcpv6 relay statistics
DHCPv6 Packets dropped:
    Total 1
    Lease Time Violation 1

Messages received:
    DHCPV6_DECLINE 0
    DHCPV6_SOLICIT 10
    DHCPV6_INFORMATION_REQUEST 0
    DHCPV6_RELEASE 0
    DHCPV6_REQUEST 10
    DHCPV6_CONFIRM 0
    DHCPV6_RENEW 0
    DHCPV6_REBIND 0
    DHCPV6_RELAY_REPL 0

Messages sent:
    DHCPV6_ADVERTISE 0
    DHCPV6_REPLY 0

```

DHCPV6_RECONFIGURE	0
DHCPV6_RELAY_FORW	0
Packets forwarded:	
Total	4
FWD REQUEST	2
FWD REPLY	2
External Server Response:	
State	Responding

## show route extensive

<b>List of Syntax</b>	<a href="#">Syntax on page 163</a> <a href="#">Syntax (EX Series Switches) on page 163</a>
<b>Syntax</b>	show route extensive <destination-prefix> <logical-system (all   logical-system-name)>
<b>Syntax (EX Series Switches)</b>	show route extensive <destination-prefix>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Display extensive information about the active entries in the routing tables.
<b>Options</b>	<b>none</b> —Display all active entries in the routing table.  <b>destination-prefix</b> —(Optional) Display active entries for the specified address or range of addresses.  <b>logical-system (all   logical-system-name)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show route extensive on page 170</a> <a href="#">show route extensive (Access Route) on page 176</a> <a href="#">show route extensive (BGP PIC Edge) on page 177</a> <a href="#">show route extensive (FRR and LFA) on page 177</a> <a href="#">show route extensive (Route Reflector) on page 178</a> <a href="#">show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 178</a> <a href="#">show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 179</a>
<b>Output Fields</b>	Table 14 on page 163 describes the output fields for the <b>show route extensive</b> command. Output fields are listed in the approximate order in which they appear.

Table 14: show route extensive Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.

Table 14: show route extensive Output Fields (*continued*)

Field Name	Field Description
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active).</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive).</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy).</li> </ul>
<i>route-destination</i> (entry, announced)	<p>Route destination (for example: 10.0.0.1/24). The <b>entry</b> value is the number of route for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> </ul>
<b>TSI</b>	Protocol header information.
<b>label stacking</b>	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the Multiprotocol Label Switching (MPLS) label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>
<b>[protocol, preference]</b>	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+—</b>A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>- —</b>A hyphen indicates the last active route.</li> <li>• <b>*—</b>An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>

Table 14: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b>Level</b>	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
<b>Route Distinguisher</b>	IP subnet augmented with a 64-bit prefix.
<b>PMSI</b>	Provider multicast service interface (MVPN routing table).
<b>Next-hop type</b>	Type of next hop. For a description of possible values for this field, see the Output Field table in the <i>show route detail</i> command.
<b>Next-hop reference count</b>	Number of references made to the next hop.
<b>Flood nexthop branches exceed maximum message</b>	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
<b>Source</b>	IP address of the route source.
<b>Next hop</b>	Network layer address of the directly reachable neighboring system.
<b>via</b>	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b>. This field can also contain the following information:</p> <ul style="list-style-type: none"> <li>• <b>Weight</b>—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>• <b>Balance</b>—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing.</li> </ul>
<b>Label-switched-path lsp-path-name</b>	Name of the label-switched path (LSP) used to reach the next hop.
<b>Label operation</b>	MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).
<b>Offset</b>	Whether the metric has been increased or decreased by an offset value.
<b>Interface</b>	(Local only) Local interface name.
<b>Protocol next hop</b>	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Table 14: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b><i>label-operation</i></b>	MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).
<b>Indirect next hops</b>	<p>When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.</p> <p>When BGP PIC Edge is enabled, the output lines that contain <b>Indirect next hop: weight</b> follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none"> <li>• 0x1 indicates active next hops.</li> <li>• 0x4000 indicates passive next hops.</li> </ul>
<b>State</b>	State of the route (a route can be in more than one state). See the Output Field table in the <i>show route detail</i> command.
<b>Session ID</b>	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
<b>Weight</b>	<p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see <b>show route table</b>.</p>

Table 14: show route extensive Output Fields (*continued*)

Field Name	Field Description
Inactive reason	<p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> <li>• <b>Active preferred</b>—Currently active route was selected over this route.</li> <li>• <b>Always compare MED</b>—Path with a lower multiple exit discriminator (MED) is available.</li> <li>• <b>AS path</b>—Shorter AS path is available.</li> <li>• <b>Cisco Non-deterministic MED selection</b>—Cisco nondeterministic MED is enabled and a path with a lower MED is available.</li> <li>• <b>Cluster list length</b>—Path with a shorter cluster list length is available.</li> <li>• <b>Forwarding use only</b>—Path is only available for forwarding purposes.</li> <li>• <b>IGP metric</b>—Path through the next hop with a lower IGP metric is available.</li> <li>• <b>IGP metric type</b>—Path with a lower OSPF link-state advertisement type is available.</li> <li>• <b>Interior &gt; Exterior &gt; Exterior via Interior</b>—Direct, static, IGP, or EBGp path is available.</li> <li>• <b>Local preference</b>—Path with a higher local preference value is available.</li> <li>• <b>Next hop address</b>—Path with a lower metric next hop is available.</li> <li>• <b>No difference</b>—Path from a neighbor with a lower IP address is available.</li> <li>• <b>Not Best in its group</b>—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed).</li> <li>• <b>Number of gateways</b>—Path with a higher number of next hops is available.</li> <li>• <b>Origin</b>—Path with a lower origin code is available.</li> <li>• <b>OSPF version</b>—Path does not support the indicated OSPF version.</li> <li>• <b>RIB preference</b>—Route from a higher-numbered routing table is available.</li> <li>• <b>Route distinguisher</b>—64-bit prefix added to IP subnets to make them unique.</li> <li>• <b>Route metric or MED comparison</b>—Route with a lower metric or MED is available.</li> <li>• <b>Route preference</b>—Route with a lower preference value is available.</li> <li>• <b>Router ID</b>—Path through a neighbor with a lower ID is available.</li> <li>• <b>Unusable path</b>—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved.</li> <li>• <b>Update source</b>—Last tiebreaker is the lowest IP address value.</li> </ul>
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signalled and LDP-signalled LSPs or for specific VRF routing instances.</p> <p>For sample output, see <b>show route table</b>.</p>

Table 14: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b>Task</b>	Name of the protocol that has added the route.
<b>Announcement bits</b>	List of protocols that announce this route. <b>n-Resolve inet</b> indicates that the route is used for route resolution for next hops found in the routing table. <b>n</b> is an index used by Juniper Networks customer support only.
<b>AS path</b>	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li>• <b>I</b>—IGP.</li> <li>• <b>E</b>—EGP.</li> <li>• <b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li>• <b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li>• <b>[ ]</b>—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured.</li> <li>• <b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li>• <b>( )</b>—Parentheses enclose a confederation.</li> <li>• <b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
<b>validation-state</b>	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>
<b>FECs bound to route</b>	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
<b>AS path: I &lt;Originator&gt;</b>	(For route reflected output only) Originator ID attribute set by the route reflector.



Table 14: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b>route status</b>	<p>Indicates the status of a BGP route:</p> <ul style="list-style-type: none"> <li>• <b>Accepted</b>—The specified BGP route is imported by the default BGP policy.</li> <li>• <b>Import</b>—The route is imported into a Layer 3 VPN routing instance.</li> <li>• <b>Import-Protect</b>—A remote instance egress that is protected.</li> <li>• <b>Multipath</b>—A BGP multipath active route.</li> <li>• <b>MultipathContrib</b>—The route is not active but contributes to the BGP multipath.</li> <li>• <b>Protect</b>—An egress route that is protected.</li> <li>• <b>Stale</b>—A route that is marked stale due to graceful restart.</li> </ul>
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the <i>show route detail</i> command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: <b>none</b> or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.

Table 14: show route extensive Output Fields (*continued*)

Field Name	Field Description
<b>status vector</b>	Layer 2 VPN and VPLS network layer reachability information (NLRI).
<b>Localpref</b>	Local preference value included in the route.
<b>Router ID</b>	BGP router ID as advertised by the neighbor in the open message.
<b>Primary Routing Table</b>	In a routing table group, the name of the primary routing table in which the route resides.
<b>Secondary Tables</b>	In a routing table group, the name of one or more secondary tables in which the route resides.
<b>Originating RIB</b>	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
<b>Node path count</b>	Number of nodes in the path.
<b>Forwarding nexthops</b>	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

## Sample Output

### show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:34:06
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:32:40
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/3/0.0, selected

```

```

        State: <Int>
        Inactive reason: Route Preference
        Local AS: 69
        Age: 1:32:40    Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

10.31.1.1/32 (1 entry, 1 announced)
    *Local Preference: 0
        Next hop type: Local
        Next-hop reference count: 7
        Interface: so-0/3/0.0
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:32:43
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kerne 10.31.2.0/30 -> {10.31.1.6}
    *OSPF Preference: 10
        Next-hop reference count: 9
        Next hop: via so-0/3/0.0
        Next hop: 10.31.1.6 via ge-3/1/0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:32:19    Metric: 2
        Area: 0.0.0.0
        Task: OSPF
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kerne 224.0.0.2/32 -> {}
    *PIM Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:08
        Task: PIM Recv
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kerne 224.0.0.22/32 -> {}
    *IGMP Preference: 0
        Next-hop reference count: 18
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:34:06

```

```

Task: IGMP
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
*RSVP Preference: 7
Next-hop reference count: 6
Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
Label-switched-path green-r1-r3
Label operation: Push 100096
State: <Active Int>
Local AS: 69
Age: 1:28:12 Metric: 2
Task: RSVP
Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
AS path: I

10.255.71.238/32 (1 entry, 1 announced)
State: <FlashAll>
*RSVP Preference: 7
Next-hop reference count: 6
Next hop: via so-0/3/0.0 weight 0x1, selected
Label-switched-path green-r1-r2
State: <Active Int>
Local AS: 69
Age: 1:28:12 Metric: 1
Task: RSVP
Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 69
Age: 1:34:07
Task: IF
AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
*MPLS Preference: 0
Next hop type: Receive
Next-hop reference count: 6
State: <Active Int>
Local AS: 69
Age: 1:34:08 Metric: 1

```

```

Task: MPLS
Announcement bits (1): 0-KRT
AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299776 (1 entry, 1 announced)
TSI:
KRT in-kernel 299776 /52 -> {Flood}
    *RSVP   Preference: 7
            Next hop type: Flood
            Next-hop reference count: 130
            Flood nexthop branches exceed maximum
            Address: 0x8ea65d0

...

800010 (1 entry, 1 announced)
TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
    *VPLS   Preference: 7
            Next-hop reference count: 2
            Next hop: via vt-3/2/0.32769, selected
            Label operation: Pop
            State: <Active Int>
            Age: 1:31:53
            Task: Common L2 VC
            Announcement bits (1): 0-KRT
            AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
    *VPLS   Preference: 7
            Next-hop reference count: 2
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
            Label-switched-path green-r1-r3
            Label operation: Push 800012, Push 100096(top)
            Protocol next hop: 10.255.70.103
            Push 800012
            Indirect next hop: 87272e4 1048574
            State: <Active Int>
            Age: 1:31:53   Metric2: 2
            Task: Common L2 VC
            Announcement bits (2): 0-KRT 1-Common L2 VC
            AS path: I
            Communities: target:11111:1 Layer2-info: encaps:VPLS,
            control flags:, mtu: 0
            Indirect next hops: 1
                Protocol next hop: 10.255.70.103 Metric: 2
                Push 800012
                Indirect next hop: 87272e4 1048574
                Indirect path forwarding next hops: 1
                    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
                    10.255.70.103/32 Originating RIB: inet.3
                    Metric: 2   Node path count: 1
                    Forwarding nexthops: 1
                    Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

```

```
abcd::10:255:71:52/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.16385, selected
    State: <Active NoReadvrt Int>
    Age: 1:34:07
    Task: IF
    AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:28:12 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0

```

```

Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  55.0.0.0/24 -> {Push 300112}
    *BGP Preference: 170/-101
      Next hop type: Router
      Address: 0x925c208
      Next-hop reference count: 2
      Source: 10.0.0.9
      Next hop: 10.0.0.9 via lt-1/2/0.15, selected
      Label operation: Push 300112
      Label TTL action: prop-ttl
      State: <Active Ext>
      Local AS: 7019 Peer AS: 13979
      Age: 1w0d 23:06:56
      AIGP: 25
      Task: BGP_13979.10.0.0.9+56732
      Announcement bits (1): 0-KRT
      AS path: 13979 7018 I
      Accepted
      Route Label: 300112
      Localpref: 100
      Router ID: 10.9.9.1

```

### show route extensive (Access Route)

```

user@host> show route 13.160.0.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
  *Access Preference: 13
    Next-hop reference count: 78472
    Next hop: 13.160.0.2 via fe-0/0/0.0, selected
    State: <Active Int>

Age: 12
Task: RPD Unix Domain Server./var/run/rpd_serv.local

```



```
Announcement bits (2): 0-KRT 1-OSPFv2
AS path: I
```

### show route extensive (BGP PIC Edge)

```
user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  1.1.1.6/32 (3 entries, 2 announced)
    State: <CalcForwarding>
    TSI:
    KRT in-kernel 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
    Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
    Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
  ..
    #Multipath Preference: 255
      Next hop type: Indirect
      Address: 0x93f4010
      Next-hop reference count: 2
  ..
    Protocol next hop: 1.1.1.4
    Push 299824
    Indirect next hop: 944c000 1048574 INH Session ID: 0x3
    Indirect next hop: weight 0x1
    Protocol next hop: 1.1.1.5
    Push 299824
    Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
    Indirect next hop: weight 0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 25      Metric2: 15
    Validation State: unverified
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: 3 I
    Communities: target:2:1
```

### show route extensive (FRR and LFA)

```
user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
  20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll
    TSI:
    KRT in-kernel 20.31.2.0/24 -> {Push 299776, Push 299792}
      *RSVP Preference: 7/1
      Next hop type: Router, Next hop index: 1048574
      Address: 0xbbbc010
      Next-hop reference count: 5
      Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
      Label-switched-path europa-d-to-europa-e
      Label operation: Push 299776
      Label TTL action: prop-ttl
      Session Id: 0x201
      Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
      Label-switched-path europa-d-to-europa-e
      Label operation: Push 299792
      Label TTL action: prop-ttl
      Session Id: 0x202
```

```

State: Active Int
Local AS: 100
Age: 5:31 Metric: 2
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

#### show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
*BGP Preference: 170/-101
Source: 192.168.4.214
Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
State: <Active Int Ext>
Local AS: 10458 Peer AS: 10458
Age: 3:09 Metric: 0 Metric2: 0
Task: BGP_10458.192.168.4.214+1033
Announcement bits (2): 0-KRT 4-Resolve inet.0
AS path: 3944 7777 I <Originator>
Cluster list: 1.1.1.1
Originator ID: 10.255.245.88
Communities: 7777:7777
Localpref: 100
Router ID: 4.4.4.4
Indirect next hops: 1
    Protocol next hop: 207.17.136.192 Metric: 0
    Indirect next hop: 84ac908 40
    Indirect path forwarding next hops: 0
    Next hop type: Discard

```

#### show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Flood
Next-hop reference count: 3
Address: 0x9097d90
Next hop: via vt-0/1/0.1
Next-hop index: 661
Label operation: Pop
Address: 0x9172130
Next hop: via so-0/0/3.0

```

```

Next-hop index: 654
Label operation: Swap 299872
State: **Active Int>
Local AS: 1001
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

### show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP   Preference: 9
        Next hop type: Flood
        Address: 0x2735208
        Next-hop reference count: 3
        Next hop type: Router, Next hop index: 1397
        Address: 0x2735d2c
        Next-hop reference count: 3
        Next hop: 1.3.8.2 via ge-1/2/22.0
        Label operation: Pop
        Load balance label: None;
        Next hop type: Router, Next hop index: 1395
        Address: 0x2736290
        Next-hop reference count: 3
        Next hop: 1.3.4.2 via ge-1/2/18.0
        Label operation: Pop
        Load balance label: None;
        State: <Active Int AckRequest MulticastRPF>
        Local AS: 10
        Age: 54:05      Metric: 1
        Validation State: unverified
        Task: LDP
        Announcement bits (1): 0-KRT
        AS path: I
        FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
        Primary Upstream : 1.1.1.3:0--1.1.1.2:0
          RPF Nexthops :
            ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
            ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
        Backup Upstream : 1.1.1.3:0--1.1.1.6:0
          RPF Nexthops :
            ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
            ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

## show route protocol

---

<b>List of Syntax</b>	<a href="#">Syntax on page 180</a> <a href="#">Syntax (EX Series Switches) on page 180</a>
<b>Syntax</b>	<code>show route protocol <i>protocol</i></code> <brief   detail   extensive   terse> <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	<code>show route protocol <i>protocol</i></code> <brief   detail   extensive   terse>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Options <b>ospf2</b> and <b>ospf3</b> introduced in Junos OS Release 9.2. Options <b>ospf2</b> and <b>ospf3</b> introduced in Junos OS Release 9.2 for EX Series switches. Option <b>flow</b> introduced in Junos OS Release 10.0. Option <b>flow</b> introduced in Junos OS Release 10.0 for EX Series switches.
<b>Description</b>	Display the route entries in the routing table that were learned from a particular protocol.
<b>Options</b>	<b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b><i>protocol</i></b> —Protocol from which the route was learned: <ul style="list-style-type: none"><li>• <b>access</b>—Access route for use by DHCP application</li><li>• <b>access-internal</b>—Access-internal route for use by DHCP application</li><li>• <b>aggregate</b>—Locally generated aggregate route</li><li>• <b>atmvpn</b>—Asynchronous Transfer Mode virtual private network</li><li>• <b>bgp</b>—Border Gateway Protocol</li><li>• <b>ccc</b>—Circuit cross-connect</li><li>• <b>direct</b>—Directly connected route</li><li>• <b>dvmrp</b>—Distance Vector Multicast Routing Protocol</li><li>• <b>esis</b>—End System-to-Intermediate System</li><li>• <b>flow</b>—Locally defined flow-specification route.</li><li>• <b>isis</b>—Intermediate System-to-Intermediate System</li><li>• <b>ldp</b>—Label Distribution Protocol</li><li>• <b>l2circuit</b>—Layer 2 circuit</li><li>• <b>l2vpn</b>—Layer 2 virtual private network</li><li>• <b>local</b>—Local address</li></ul>

- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First version 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



**NOTE:** EX Series switches run a subset of these protocols. See the switch CLI for details.

<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show route protocol access on page 182</a> <a href="#">show route protocol access-internal extensive on page 182</a> <a href="#">show route protocol bgp on page 182</a> <a href="#">show route protocol bgp detail on page 182</a> <a href="#">show route protocol bgp extensive on page 183</a> <a href="#">show route protocol bgp terse on page 183</a> <a href="#">show route protocol direct on page 183</a> <a href="#">show route protocol l2circuit detail on page 184</a> <a href="#">show route protocol l2vpn extensive on page 185</a> <a href="#">show route protocol ldp on page 185</a> <a href="#">show route protocol ldp extensive on page 186</a> <a href="#">show route protocol ospf (Layer 3 VPN) on page 187</a> <a href="#">show route protocol ospf detail on page 188</a> <a href="#">show route protocol rip on page 188</a> <a href="#">show route protocol rip detail on page 188</a> <a href="#">show route protocol ripng table inet6 on page 188</a>
<b>Output Fields</b>	For information about output fields, see the output field tables for the <b>show route</b> command, the <b>show route detail</b> command, the <b>show route extensive</b> command, or the <b>show route terse</b> command.

## Sample Output

### show route protocol access

```

user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0

```

### show route protocol access-internal extensive

```

user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
      Next-hop reference count: 200000
      Next hop: 13.160.0.2 via fe-0/0/0.0, selected
      State: <Active Int>
    Age: 36
      Task: RPD Unix Domain Server./var/run/rpd_serv.local
      Announcement bits (1): 0-KRT
      AS path: I

```

### show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21    *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                  AS path: 10458 14203 2914 4788 4788 I
                  > to 192.168.167.254 via fxp0.0

```

### show route protocol bgp detail

```

show route protocol bgp 66.117.63.0/24 exact detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
      Next hop type: Indirect
      Next-hop reference count: 1006436
      Source: 192.168.69.71
      Next hop type: Router, Next hop index: 324
      Next hop: 192.168.167.254 via fxp0.0, selected
      Protocol next hop: 192.168.69.71
      Indirect next hop: 8e166c0 342
      State: <Active Ext>
      Local AS: 69 Peer AS: 10458
      Age: 6d 10:42:42      Metric2: 0
      Task: BGP_10458.192.168.69.71+179
      Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
      AS path: 10458 14203 2914 4788 4788 I

```

```

Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192

```

### show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive
```

```
inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
```

```

KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
  *BGP Preference: 170/-101
    Next hop type: Indirect
    Next-hop reference count: 1006502
    Source: 192.168.69.71
    Next hop type: Router, Next hop index: 324
    Next hop: 192.168.167.254 via fxp0.0, selected
    Protocol next hop: 192.168.69.71
    Indirect next hop: 8e166c0 342
    State: <Active Ext>
    Local AS: 69 Peer AS: 10458
    Age: 6d 10:44:45 Metric2: 0
    Task: BGP_10458.192.168.69.71+179
    Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

```

```
1
```

```

AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
  Protocol next hop: 192.168.69.71
  Indirect next hop: 8e166c0 342
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 192.168.167.254 via fxp0.0
192.168.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 192.168.167.254 via fxp0.0

```

### show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse
```

```
inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
192.168.64.0/21	B 170	100		>100.1.3.2	10023 21 I

### show route protocol direct

```
user@host> show route protocol direct
```

```

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24          *[Direct/0] 17w0d 10:31:49
                   > via fe-1/3/1.0
10.255.165.1/32    *[Direct/0] 25w4d 04:13:18
                   > via lo0.0
30.30.30.0/24      *[Direct/0] 17w0d 23:06:26
                   > via fe-1/3/2.0
192.168.164.0/22   *[Direct/0] 25w4d 04:13:20
                   > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                   *[Direct/0] 25w4d 04:13:21
                   > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

abcd::10:255:165:1/128
                   *[Direct/0] 25w4d 04:13:21
                   > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                   *[Direct/0] 25w4d 04:13:21
                   > via lo0.0

```

### show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
   *L2CKT Preference: 7
       Next hop: via ge-2/0/0.0, selected
       Label operation: Pop          Offset: 4
       State: <Active Int>
       Local AS: 99
       Age: 9:52
       Task: Common L2 VC
       Announcement bits (1): 0-KRT
       AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
   *L2CKT Preference: 7
       Next hop: via so-1/1/2.0 weight 1, selected
       Label-switched-path my-lsp
       Label operation: Push 100000, Push 100000(top)[0] Offset: -4
       Protocol next hop: 10.245.255.63
       Push 100000 Offset: -4
       Indirect next hop: 86af0c0 298
       State: <Active Int>
       Local AS: 99
       Age: 9:52
       Task: Common L2 VC
       Announcement bits (2): 0-KRT 1-Common L2 VC
       AS path: I

```



```

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

### show route protocol l2vpn extensive

```

user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected
    Label operation: Push 800000 Offset: -4
    Protocol next hop: 10.255.14.220
    Push 800000 Offset: -4
    Indirect next hop: 85142a0 288
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:69:1 Layer2-info: encaps:PPP,
    control flags:2, mtu: 0

```

### show route protocol ldp

```

user@host> show route protocol ldp

```

```

inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000

```

### show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
  *LDP    Preference: 9
          Next-hop reference count: 2
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>

```

```

Local AS: 65500
Age: 1d 23:03:58      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
  *LDP      Preference: 9
            Next-hop reference count: 2
            Next hop: via t1-4/0/0.0, selected
            Label operation: Pop
            State: <Active Int>
            Local AS: 65500
            Age: 1d 23:03:58      Metric: 1
            Task: LDP
            Announcement bits (1): 0-KRT
            AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
  *LDP      Preference: 9
            Next-hop reference count: 2
            Next hop: via t1-4/0/0.0, selected
            Label operation: Swap 100000
            State: <Active Int>
            Local AS: 65500
            Age: 1d 23:03:58      Metric: 1
            Task: LDP
            Announcement bits (1): 0-KRT
            AS path: I
            Prefixes bound to route: 192.168.16.1/32

```

### show route protocol ospf (Layer 3 VPN)

```

user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.255.14.179/32  *[OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
224.0.0.5/32      *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
10.255.14.173/32  *[OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
224.0.0.5/32      *[OSPF/10] 20:26:20, metric 1

```

### show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
    OSPF   Preference: 10
           Nexthop: via so-0/2/2.0, selected
           State: <Int>
           Inactive reason: Route Preference
           Age: 6:25      Metric: 1
           Area: 0.0.0.0
           Task: VPN-AB-OSPF
           AS path: I
           Communities: Route-Type:0.0.0.0:1:0

...

```

### show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  * [RIP/100] 20:24:34, metric 2
                  > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32      * [RIP/100] 00:03:59, metric 1

```

### show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
    *RIP   Preference: 100
           Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
           State: <Active Int>
           Age: 20:25:02  Metric: 2
           Task: VPN-AB-RIPv2
           Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
           AS path: I
           Route learned from 10.39.1.22 expires in 96 seconds

```

### show route protocol ripng table inet6

```

user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      * [RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0

```

```
1111::4/128      *[RIPng/100] 02:13:33, metric 2  
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0  
1111::5/128      *[RIPng/100] 02:13:33, metric 2  
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0  
1111::6/128      *[RIPng/100] 02:13:33, metric 2  
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

