



Services Feature Guide for the QFX Series

Release

14.1X53



Modified: 2015-11-25

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

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

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

Services Feature Guide for the QFX Series
14.1X53
Copyright © 2015, Juniper Networks, Inc.
All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xi
	Documentation and Release Notes	xi
	Supported Platforms	xi
	Using the Examples in This Manual	xi
	Merging a Full Example	xii
	Merging a Snippet	xii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xv
	Opening a Case with JTAC	xvi
Part 1	Port Mirroring	
Chapter 1	Using Port Mirroring	3
	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
	Understanding Layer 3 Logical Interfaces	8
	Examples: Configuring Port Mirroring for Local Analysis	9
	Example: Mirroring Employee Web Traffic with a Firewall Filter	11
	Example: Configuring Port Mirroring for Remote Analysis	14
	Example: Mirroring Employee Web Traffic with a Firewall Filter	19
	Configuring Port Mirroring	22
	Configuring Port Mirroring for Local Analysis	23
	Configuring Port Mirroring for Remote Analysis	23
	Filtering the Traffic Entering an Analyzer	24
	Troubleshooting Port Mirroring	25
	Port Mirroring Constraints and Limitations	25
	Local and Remote Port Mirroring	25
	Remote Port Mirroring Only	27
	Egress Port Mirroring with VLAN Translation	27
	Egress Port Mirroring with Private VLANs	27

Part 2	DHCP and DHCP Relay	
Chapter 2	Using DHCP and DHCP Relay	31
	DHCP and BOOTP Relay Overview	31
	Configuring DHCP and BOOTP	32
	Configuring DHCP Relay and Smart Relay	32
	Configuring DHCP Relay	32
	Configuring DHCP Smart Relay	32
Part 3	Configuration Statements and Operational Commands	
Chapter 3	Configuration Statements for Port Mirroring	37
	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
Chapter 4	Configuration Statements for DHCP and DHCP Relay	53
	apply-secondary-as-giaddr	54
	dhcp-local-server	55
	dhcp-relay	60
Chapter 5	Configuration Statements for Encryption	67
	authentication-key-chains	68
	cache-size	69
	cache-timeout-negative	70
	ca-name	71
	certificates	72
	certification-authority	73
	crl (Encryption Interface)	74
	encoding	74
	enrollment-retry	75
	enrollment-url	75
	file	76
	key (Authentication Keychain)	77
	key-chain (Security)	78
	ldap-url	79
	local	80
	maximum-certificates	81
	path-length	81

	secret	82
	security	83
	ssh-known-hosts	84
	start-time (Authentication Key Transmission)	85
	traceoptions	87
Chapter 6	Operational Command for Port Mirroring	89
	show analyzer	90

List of Figures

Part 1	Port Mirroring	
Chapter 1	Using Port Mirroring	3
	Figure 1: Network Topology for Local Port Mirroring Example	10

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xiii
	Table 2: Text and Syntax Conventions	xiii
Part 1	Port Mirroring	
Chapter 1	Using Port Mirroring	3
	Table 3: Port Mirroring Terms and Definitions	4
Part 3	Configuration Statements and Operational Commands	
Chapter 6	Operational Command for Port Mirroring	89
	Table 4: show analyzer Output Fields	90

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
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

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

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

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

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

Documentation Feedback

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

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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active 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

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

CHAPTER 1

Using Port Mirroring

- [Understanding Port Mirroring on page 3](#)
- [Understanding Layer 3 Logical Interfaces on page 8](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 14](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 19](#)
- [Configuring Port Mirroring on page 22](#)
- [Troubleshooting Port Mirroring on page 25](#)

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
------	-------------

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

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).
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 port-mirror-instance <i>instance-name</i> 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"> • Cannot also be a source port. • Cannot be used for switching. • Cannot be an aggregated Ethernet interface (LAG). • Does not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP). • Loses any existing VLAN associations when you configure it as an analyzer output interface. <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"> • An output IP 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).
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"> • Cannot be a private VLAN or VLAN range. • 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 (LAG). • On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.
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 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 analyzer <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 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.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

Related Documentation

- *Configuring Port Mirroring*
- *Example: Configuring Port Mirroring for Local Analysis*
- [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 14](#)
- [Troubleshooting Port Mirroring on page 25](#)

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. .



NOTE: You can also use Layer 3 logical interfaces to provide alternative gateway addresses for smart DHCP relay

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](#)
- [Configuring DHCP and BOOTP Relay](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

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. If your switch runs software that does not support ELS, see *Example: Configuring Port Mirroring for Local Analysis*. 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 9](#)
- [Overview and Topology on page 9](#)
- [Example: Mirroring All Employee Traffic for Local Analysis on page 10](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 11](#)

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.

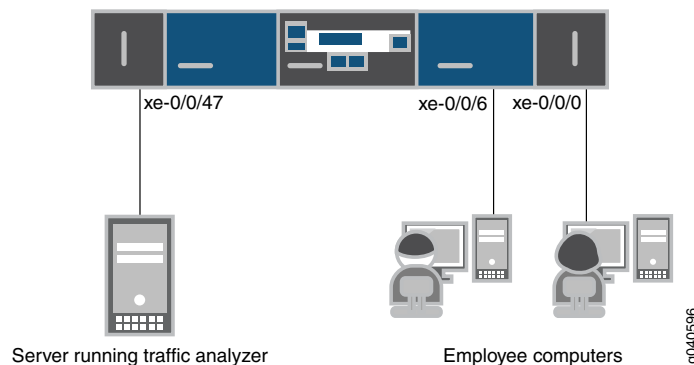
In this example, **xe-0/0/0** and **xe-0/0/6** serve as connections for employee computers. Interface **xe-0/0/47** is connected to a device running an analyzer application.



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

Figure 1 on page 10 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 11](#)
- [Overview on page 11](#)
- [Configuring on page 11](#)
- [Verification on page 13](#)

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	<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>
Step-by-Step Procedure	<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"> 1. Configure the output interface: <pre>[edit interfaces] user@switch# set xe-0/0/47 unit 0 family ethernet-switching</pre> 2. Configure the employee-web-monitor 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> 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: <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> 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>
Results	<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](#)

Related Documentation

- [Understanding Port Mirroring on page 3](#)
- [Configuring Port Mirroring on page 22](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 14](#)

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.



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

-
- [Requirements on page 15](#)
 - [Overview and Topology on page 15](#)
 - [Mirroring All Employee Traffic for Remote Analysis on page 15](#)

- [Mirroring Employee-to-Web Traffic for Remote Analysis on page 16](#)
- [Verification on page 18](#)

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 port mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set forwarding-options analyzer 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 analyzer
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 port 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 analyzer 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 analyzer
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 {
        port-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 analyzer employee-web-monitor;
      }
    }
  }
}
forwarding-options analyzer {
  employee-web-monitor {
    output {
      vlan {
        999;
      }
    }
  }
}
vpls {
  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 22](#)
 - [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)
 - [Overview of Firewall Filters](#)

Example: Mirroring Employee Web Traffic with a Firewall Filter

- [Requirements on page 19](#)
- [Overview on page 19](#)
- [Configuring on page 19](#)
- [Verification on page 21](#)

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](#)

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 23](#)
- [Configuring Port Mirroring for Remote Analysis on page 23](#)
- [Filtering the Traffic Entering an Analyzer on page 24](#)

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 9](#)
- [Example: Configuring Port Mirroring for Remote Analysis on page 14](#)
- [Overview of Firewall Filters](#)

Troubleshooting Port Mirroring

- [Port Mirroring Constraints and Limitations on page 25](#)
- [Egress Port Mirroring with VLAN Translation on page 27](#)
- [Egress Port Mirroring with Private VLANs on page 27](#)

Port Mirroring Constraints and Limitations

- [Local and Remote Port Mirroring on page 25](#)
- [Remote Port Mirroring Only on page 27](#)

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.
- On the source (monitored) switch, only one interface can be a member of the analyzer VLAN.

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](#)
 - *Example: Configuring Port Mirroring for Local Analysis*
 - *Example: Configuring Port Mirroring for Remote Analysis*

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 32](#)
- [Configuring DHCP Relay and Smart Relay on page 32](#)

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.

You can configure the switch to use the gateway IP address (*giaddr*) as the source IP address of the switch for relayed DHCP packets when the switch is used as the DHCP relay agent. For information on configuring this option, see the *source-address-giaddr* configuration statement.

You can also use smart DHCP relay, which enables you to configure alternative IP addresses for the gateway interface so that if the server fails to reply to the requests sent from the primary gateway address, the switch can resend the requests using the alternative gateway addresses. To use this feature, you must configure a Layer 3 interface, Layer 3 subinterface, or IRB interface with multiple IP addresses and configure that interface to be a relay agent.



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 on page 32](#)
- [Configuring DHCP Relay on page 32](#)

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
```

Configuring DHCP Relay and Smart Relay

Configuring DHCP Relay

You can configure the QFX Series to act as a Dynamic Host Configuration Protocol (DHCP) relay agent. This means that if a locally attached host can issue a DHCP request as a broadcast message and the switch relays the message to a specified DHCP server. You should configure a switch to be a DHCP and relay agent if you have locally attached hosts and a remote DHCP or server.

To use DHCP relay, configure the **set forwarding-options dhcp-relay** statement.

Configuring DHCP Smart Relay

You can use DHCP smart relay to provide redundancy and resiliency to your DHCP relay configuration. Smart DHCP relay enables you to configure alternative IP addresses for the gateway interface so that if the server fails to reply to the requests sent from the primary gateway address, the switch can resend the requests using the alternative gateway addresses. To use DHCP smart relay, you also need an interface with multiple IP addresses assigned to it. You can achieve this by doing either of the following tasks:

- Create a Layer 3 interface and assign multiple addresses to it.

- Create a Layer 3 logical interface (by using VLAN tagging) and assign at least two IP addresses to it. See [“Understanding Layer 3 Logical Interfaces” on page 8](#) and [Configuring a Layer 3 Logical Interface](#) for information about this approach.
- Create a routed VLAN interface and assign at least two IP addresses to it. See [Configuring IRB Interfaces](#) and [Example: Configuring Routing Between VLANs on One Switch](#) for information about this approach.

When smart relay is configured for an interface, the switch will initially relay DHCP DISCOVER messages using the primary address of the interface as the gateway IP address (in the giaddr field) for the DHCP message. If no DHCP OFFER is received from a server in reply, the switch allows the client to send as many as two more DISCOVER messages and relays them using the same gateway IP address. If no DHCP OFFER is received after three DISCOVER messages, the switch starts relaying subsequent DISCOVER messages sent by client using the secondary IP address as the gateway IP address. If you configure more than two IP addresses on the relay agent interface, the switch repeats this process until a successful DHCP binding happens or all IP addresses have been used. Once all the IP addresses are used without a successful DHCP binding, the switch starts over from the primary address and the process repeats.

Enter the following statement to configure the primary interface address that smart relay will use: **set interfaces *interface* unit *unit* family inet address *address* primary**. If you do not configure a primary address, the lowest address of the multiple addresses is used.

Once you have created an interface with multiple IP addresses, complete the smart relay configuration by entering one of the following statements:

- **set forwarding-options dhcp-relay overrides apply-secondary-as-giaddr**: Use this statement to enable smart relay on all the interfaces that are configured as relay agents.
- **set forwarding-options dhcp-relay group *name* overrides apply-secondary-as-giaddr**: Use this statement to enable smart relay on a group of interfaces.
- **set forwarding-options dhcp-relay group *name* interface *name* overrides apply-secondary-as-giaddr**: Use this statement to enable smart relay on a specific interface.

PART 3

Configuration Statements and Operational Commands

- [Configuration Statements for Port Mirroring on page 37](#)
- [Configuration Statements for DHCP and DHCP Relay on page 53](#)
- [Configuration Statements for Encryption on page 67](#)
- [Operational Command for Port Mirroring on page 89](#)

CHAPTER 3

Configuration Statements for 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

```
Syntax analyzer {
    name {
        input {
            egress {
                interface (all | interface-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);
        }
    }
}
```

Hierarchy Level For platforms without ELS:

[edit ethernet-switching-options]

For platforms with ELS:

[edit forwarding-options]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
 Option **output vlan** added in Junos OS Release 12.1 for the QFX Series.
 Option **output ip-address** added in Junos OS Release 12.3 for the QFX Series for non-ELS platforms and added in 14.1X53-D10 for ELS platforms.

Description Configure port mirroring. You can create a total of four port-mirroring configurations on the QFX Series, subject to the following limits:

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

Default Port mirroring is disabled, and Junos OS creates no default analyzers.

Options **all**—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 22](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)

egress

Syntax egress {
 interface (all | *interface-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 interfaces for which egressing traffic is mirrored.

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 22](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)

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 22](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)

family (Port Mirroring)

Syntax	<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>
Hierarchy Level	[edit forwarding-options port-mirroring [<i>instance name</i>]]
Release Information	Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	Specify the type of interface that will be used to forward port mirrored packet to an analyzer device..
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Port Mirroring on page 3 • Configuring Port Mirroring on page 22 • Examples: Configuring Port Mirroring for Local Analysis on page 9

inet (Port Mirroring)

Syntax	<pre> inet { output { ip-address address { } routing-instance instance-name { ip-address address { } } } } </pre>
Hierarchy Level	[edit forwarding-options port-mirroring [instance <i>name</i>] family]
Release Information	<p>Statement introduced in Junos OS Release 14.1X53 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>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.)</p>
	<p> NOTE: An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</p>
	<p> 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).</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Port Mirroring</i> • <i>Example: Mirroring Employee Web Traffic with a Firewall Filter</i> • <i>Configuring Port Mirroring</i>

ingress (Port Mirroring)

Syntax	<pre>ingress { interface (all interface-name); vlan (vlan-id vlan-name); }</pre>
Hierarchy Level	<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>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<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>
Required Privilege Level	<p>routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Port Mirroring on page 3 • Configuring Port Mirroring on page 22 • Examples: Configuring Port Mirroring for Local Analysis on page 9

input

Syntax	<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>
Hierarchy Level	<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>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	<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>
Default	No default.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding Port Mirroring on page 3• Configuring Port Mirroring on page 22• Examples: Configuring Port Mirroring for Local Analysis on page 9

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 22](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 11](#)

interface (Port Mirroring)

Syntax	interface (all <i>interface-name</i>);
Hierarchy Level	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <i>analyzer name</i> input (egress ingress)], [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)] [edit forwarding-options <i>analyzer name</i> output] [edit forwarding-options port-mirroring[instance <i>name</i>] family ethernet-switching output]</pre>
Release Information	Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify the interfaces for which ingress traffic is mirrored. Specify the interface that mirrored traffic should be copied to (the output interface).
Options	<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 style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>CAUTION: Configuring all 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> </div>	
<p><i>interface-name</i>—Apply port mirroring to the specified interface only.</p>	
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Port Mirroring on page 3 • Configuring Port Mirroring on page 22 • Examples: Configuring Port Mirroring for Local Analysis on page 9

ip-address (Port Mirroring)

Syntax	<code>ip-address <i>ip-address</i>;</code>
Hierarchy Level	<code>[edit forwarding-options] analyzer name output]</code> <code>[edit forwarding-options port-mirroring [instance <i>name</i>] family ethernet-switching output interface name]</code>
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 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>NOTE: An output IP address cannot be in the same subnetwork as any of the switch's management interfaces.</p> </div>	
<div>  <p>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).</p> </div>	
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

output

Syntax	<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>
Hierarchy Level	<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>
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option output vlan 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>
Description	<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>
Required Privilege Level	<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 22](#)
- [Examples: Configuring Port Mirroring for Local Analysis on page 9](#)
- [Example: Mirroring Employee Web Traffic with a Firewall Filter on page 11](#)

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)

Syntax	<code>vlan (<i>vlan-id</i> <i>vlan-name</i>);</code>
Hierarchy Level	<p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options <i>analyzer name</i> input ingress], [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)] [edit forwarding-options <i>analyzer name</i> output] [edit forwarding-options port-mirroring[instance <i>name</i>] family ethernet-switching output]</pre>
Release Information	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Option output vlan added in Junos OS Release 12.1 for the QFX Series.</p>
Description	Specify that traffic entering into a VLAN should be mirrored. Configure mirrored traffic to be sent to a VLAN for remote monitoring (output).
Options	<p><i>vlan-id</i>—Numeric VLAN identifier.</p> <p><i>vlan-name</i>—Name of the VLAN.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding Port Mirroring on page 3 • Configuring Port Mirroring on page 22 • Examples: Configuring Port Mirroring for Local Analysis on page 9

CHAPTER 4

Configuration Statements for DHCP and DHCP Relay

- [apply-secondary-as-giaddr](#) on page 54
- [dhcp-local-server](#) on page 55
- [dhcp-relay](#) on page 60

apply-secondary-as-giaddr

Syntax	<code>apply-secondary-as-giaddr;</code>
Hierarchy Level	<p>For platforms without ELS:</p> <p>[edit forwarding-options helpers bootp], [edit forwarding-options helpers bootp interface]</p> <p>For platforms with ELS:</p> <p>[edit forwarding-options dhcp-relay overrides] [edit forwarding-options dhcp-relay group <i>name</i> interface <i>name</i> overrides]</p>
Release Information	Statement introduced in Junos OS Release 12.3 for QFX Series switches.
Description	<p>Configures the interfaces on a switch that are DHCP relay agents to be enabled for smart DHCP relay. Smart DHCP relay enables you to configure alternative IP addresses for the gateway interface so that if the server fails to reply to the requests sent from the primary gateway address, the switch can resend the requests using the alternative gateway addresses. To use this feature, you must configure an IRB interface or Layer 3 subinterface with multiple IP addresses and configure that interface to be a relay agent .</p> <p>Depending on where you configure this statement, it enables smart relay on all the interfaces that are relay agents or on specific interfaces. To enable smart relay on all the interfaces that are relay agents, configure this statement at the following locations:</p> <ul style="list-style-type: none">• Junos OS with ELS: configure it under edit forwarding-options dhcp-relay overrides.• Junos OS without ELS: configure it directly under the bootp statement.. <p>To enable smart relay on specific interfaces that are relay agents, configure this statement at the following locations:</p> <ul style="list-style-type: none">• Junos OS with ELS: configure it under edit forwarding-options dhcp-relay group <i>name</i> interface <i>name</i> overrides.• Junos OS without ELS: configure it directly under the interface statement..
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring DHCP and BOOTP Relay</i>

dhcp-local-server

```
Syntax  dhcp-local-server {
        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 {
            authentication {
                ...
            }
            group group-name {
                authentication {
                    ...
                }
                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 {
                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;
  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);
    include-option-82 {
        forcerenew;
        nak;
    }
    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 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 **dhcpxv6** 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*
- *Configuring a DHCP Server on Switches (CLI Procedure)*

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 {
    ...
}
relay-option {
    ...
}
service-profile dynamic-profile-name;
trace;
upto upto-interface-name;
}
route-suppression:
service-profile dynamic-profile-name;
overrides {
    ...
}
relay-agent-interface-id {
    ...
}
relay-agent-remote-id {
    ...
}
relay-option {
    ...
}
route-suppression;
server-response-time seconds;
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;
    delay-authentication;
    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);
    use-option-82;
}
relay-agent-remote-id {
    prefix prefix;
    use-interface-description (logical | device);
}
server-group {
    server-group-name {
        server-ip-address;
    }
}
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 {
    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 {
    ...
}
relay-option-82 {
    ...
}
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 {

```

```
allow-snooped-clients;
always-write-giaddr;
always-write-option-82;
client-discover-match (option60-and-option82 | incoming-interface);
delay-authentication;
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 {
  option-number option-number;
  default-action {
    drop;
    forward-only;
    relay-server-group group-name;
  }
  equals (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    relay-server-group relay-server-group;
  }
  starts-with (ascii ascii-string | hexadecimal hexadecimal-string) {
    drop;
    forward-only;
    local-server-group local-server-group;
    relay-server-group relay-server-group;
  }
}
}
relay-option-82 {
  circuit-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
  remote-id {
    prefix prefix;
    use-interface-description (logical | device);
  }
}
}
server-group {
  server-group-name {
    server-ip-address;
  }
}
}
route-suppression:
server-response-time seconds;
service-profile dynamic-profile-name;
}
```


Hierarchy Level	<p>[edit forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> forwarding-options],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> forwarding-options],</p> <p>[edit routing-instances <i>routing-instance-name</i> forwarding-options]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	<p>Configure extended Dynamic Host Configuration Protocol (DHCP) relay and DHCPv6 relay options on the router or switch and enable the router (or 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.</p> <p>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 or client authentication. You can attach dynamic profiles and configure authentication support on a global basis or for a specific group of interfaces.</p> <p>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 (or switch) at the same time.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Extended DHCP Relay Agent Overview</i> • <i>DHCPv6 Relay Agent Overview</i> • <i>DHCP Relay Proxy Overview</i> • <i>Using External AAA Authentication Services with DHCP</i>

CHAPTER 5

Configuration Statements for Encryption

- [authentication-key-chains](#) on page 68
- [cache-size](#) on page 69
- [cache-timeout-negative](#) on page 70
- [ca-name](#) on page 71
- [certificates](#) on page 72
- [certification-authority](#) on page 73
- [crl \(Encryption Interface\)](#) on page 74
- [encoding](#) on page 74
- [enrollment-retry](#) on page 75
- [enrollment-url](#) on page 75
- [file](#) on page 76
- [key \(Authentication Keychain\)](#) on page 77
- [key-chain \(Security\)](#) on page 78
- [ldap-url](#) on page 79
- [local](#) on page 80
- [maximum-certificates](#) on page 81
- [path-length](#) on page 81
- [secret](#) on page 82
- [security](#) on page 83
- [ssh-known-hosts](#) on page 84
- [start-time \(Authentication Key Transmission\)](#) on page 85
- [traceoptions](#) on page 87

authentication-key-chains

Syntax	<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>
Hierarchy Level	[edit security]
Release Information	<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>
Description	<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 authentication-key-chains statement is configured at the [edit security] hierarchy level, and is associated with the BGP, LDP, or IS-IS protocols at the [edit protocols] hierarchy level or with the BFD protocol using the bfd-liveness-detection 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>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i> • <i>Example: Configuring BFD Authentication for Static Routes</i> • <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i>

cache-size

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



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

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

cache-timeout-negative

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



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

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

ca-name

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

certificates

Syntax	<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>
Hierarchy Level	[edit security]
Release Information	<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>
Description	<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>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Digital Certificates for an ES PIC</i>

certification-authority

Syntax	<pre>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; }</pre>
Hierarchy Level	[edit security certificates]
Release Information	<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>
Description	<p>(Encryption interface on M Series and T Series routers and EX Series switches only)</p> <p>Configure a certificate authority profile name.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Digital Certificates for an ES PIC</i>

crl (Encryption Interface)

Syntax	<code>crl file-name;</code>
Hierarchy Level	[edit security certificates]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	(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.
Options	file-name —Specify the file from which to read the CRL.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Digital Certificates for an ES PIC</i>

encoding

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

enrollment-retry

Syntax	<code>enrollment-retry <i>attempts</i>;</code>
Hierarchy Level	[edit security certificates]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	(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.
Options	<i>attempts</i> —Number of enrollment retries. Range: 0 through 100 Default: 0
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Digital Certificates for an ES PIC</i>

enrollment-url

Syntax	<code>enrollment-url <i>url-name</i>;</code>
Hierarchy Level	[edit security certificates certification-authority <i>ca-profile-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	(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).
Options	<i>url-name</i> —Certificate authority URL.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Digital Certificates for an ES PIC</i>

file

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

key (Authentication Keychain)

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


key-chain (Security)

Syntax	<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>
Hierarchy Level	[edit security authentication-key-chains]
Release Information	<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>
Description	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.
Options	<i>key-chain-name</i> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• authentication-key-chains on page 68• <i>Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols</i>• <i>Example: Configuring BFD Authentication for Static Routes</i>• <i>Example: Configuring Hitless Authentication Key Rollover for IS-IS</i>

ldap-url

Syntax	<ldap-url <i>url-name</i> >;
Hierarchy Level	[edit security certificates certification-authority <i>ca-profile-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series,
Description	(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.
Options	<i>url-name</i> —Name of the LDAP URL.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Digital Certificates for an ES PIC</i>

local

Syntax	<pre>local <i>certificate-name</i> { <i>certificate-key-string</i>; load-key-file <i>URL filename</i>; }</pre>
Hierarchy Level	[edit security certificates]
Release Information	<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>
Description	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>NOTE: 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>	
Options	<p><i>certificate-name</i><i>certificate-key-string</i>—String of alphanumeric characters that constitute the private key and certificate.</p> <p><i>certificate-name</i>—Name that uniquely identifies the certificate.</p> <p><i>load-key-file URL filename</i>—File that contains the private key and certificate. It can be one of two types of values:</p> <ul style="list-style-type: none"> • 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) • URL to the certificate file location (for instance, on the computer where the Junos XML protocol client application runs)
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Importing SSL Certificates for Junos XML Protocol Support</i>

maximum-certificates

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

path-length

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

secret

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

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)

Syntax	<code>start-time (now yyyy-mm-dd.hh:mm:ss);</code>
Hierarchy Level	[edit security authentication-key-chains key-chain <i>key-chain-name</i> key <i>key</i>]
Release Information	<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>
Description	<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>
Options	<p>now—Start time as the current year, month, day, hour, minute, and second.</p> <p>daydays—Start time as the specified number of days after the current day. For example, if the current day is the 12th and you configure start-time 2day, the start time will be on the 14th, exactly two days after the configuration is entered.</p> <p>hourhours—Start time as the specified number of hours after the current hour. For example, if the current hour is 9:00 and you configure start-time 3hour, the start time will be in 12:00, exactly three hours after the configuration is entered.</p> <p>minuteminutes—Start time as the specified number of minutes after the current minute. For example, if the current minute is 27 minutes after the hour and you configure start-time 5min, the start time will be in 32 minutes after the hour, exactly five minutes after the configuration is entered.</p> <p>monthmonths—Start time as the specified number of months after the current month. For example, if the current month is March and you configure start-time 4month, the start time will be in July, exactly four months after the configuration is entered.</p> <p>secondseconds—Start time as the specified number of seconds after the current second. For example, if the current second is 10:20:40 and you configure start-time 10seconds, the start time will be 10:20:50, exactly 10 seconds after the configuration is entered.</p> <p>yearyears—Start time as the specified number of years after the current year. For example, if the current year is 2011 and you configure start-time 1year, the start time will be in 2012, exactly one year after the configuration is entered.</p> <p>yyyy-mm-dd.hh:mm:ss—Start time in UTC (Coordinated Universal Time). The start time must be unique within the keychain.</p>
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*
 - *Example: Configuring BFD Authentication for Static Routes*
 - *Example: Configuring Hitless Authentication Key Rollover for IS-IS*

traceoptions

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>>; 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>
Hierarchy Level	<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>
Release Information	<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>
Description	<p>Configure security trace options.</p> <p>To specify more than one trace option, include multiple flag statements. Trace option output is recorded in the <code>/var/log/kmd</code> file.</p>
<div style="display: flex; align-items: center;">  <p>NOTE: The <code>traceoptions</code> statement is not supported on QFabric systems.</p> </div>	
Options	<p>files <i>number</i>—(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 size option.</p> <p>Range: 2 through 1000 files</p> <p>Default: 0 files</p> <p>size <i>size</i>—(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>Default: 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

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

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

CHAPTER 6

Operational Command for Port Mirroring

- `show analyzer`

show analyzer

Syntax	show analyzer < <i>analyzer-name</i> >
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about port mirroring.
Options	<i>analyzer-name</i> —(Optional) Displays the status of a specific analyzer (port-mirroring configuration).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Layer 2 Port Mirroring Overview</i> • Port Mirroring Constraints and Limitations on page 6 • <i>Example: Configuring Port Mirroring for Local Analysis</i> • <i>Example: Configuring Port Mirroring for Remote Analysis</i>
List of Sample Output	show analyzer on page 90
Output Fields	Table 4 on page 90 describes the output fields for the show analyzer 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.
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
```

