

# Interfaces Feature Guide for the OCX Series

Release  
14.1X53



---

Modified: 2016-08-08

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

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

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

*Interfaces Feature Guide for the OCX Series*  
14.1X53  
Copyright © 2016, Juniper Networks, Inc.  
All rights reserved.

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

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

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

# Table of Contents

	About the Documentation . . . . .	xi
	Documentation and Release Notes . . . . .	xi
	Supported Platforms . . . . .	xi
	Using the Examples in This Manual . . . . .	xi
	Merging a Full Example . . . . .	xii
	Merging a Snippet . . . . .	xii
	Documentation Conventions . . . . .	xiii
	Documentation Feedback . . . . .	xv
	Requesting Technical Support . . . . .	xv
	Self-Help Online Tools and Resources . . . . .	xv
	Opening a Case with JTAC . . . . .	xvi
<b>Part 1</b>	<b>Interfaces</b>	
<b>Chapter 1</b>	<b>Understanding Interfaces . . . . .</b>	<b>3</b>
	Interfaces Overview . . . . .	3
	Network Interfaces . . . . .	3
	Special Interfaces . . . . .	4
	Understanding Interface Naming Conventions . . . . .	4
	Physical Part of an Interface Name . . . . .	4
	Wildcard Characters in Interface Names . . . . .	5
	Understanding Interface Ranges . . . . .	5
	Understanding Management Interfaces . . . . .	6
	Configuring the Interface Address . . . . .	8
	Configuring Interface IPv4 Addresses . . . . .	9
	Operational Behavior of Interfaces When the Same IPv4 Address Is Assigned to Them . . . . .	10
	Configuring Interface IPv6 Addresses . . . . .	12
	Configuring Gigabit and 10-Gigabit Ethernet Interfaces . . . . .	13
	Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces . . . . .	13
	Configuring the IP Options . . . . .	13
	Configuring Ethernet Loopback Capability . . . . .	14
	Configuring an LPM Table With Junos OS Release 13.2X51-D10 . . . . .	15
	Monitoring Interface Status and Traffic . . . . .	17
	Troubleshooting Network Interfaces . . . . .	17
	The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down . . . . .	17

<b>Part 2</b>	<b>Generic Routing Encapsulation (GRE)</b>	
<b>Chapter 2</b>	<b>Understanding GRE</b>	<b>21</b>
	Understanding Generic Routing Encapsulation	21
	Overview of GRE	21
	GRE Tunneling	22
	Encapsulation and De-Encapsulation on the Switch	22
	Number of Source and Destination Tunnels Allowed on a Switch	22
	Class of Service on GRE Tunnels	23
	Applying Firewall Filters to GRE Traffic	23
	Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches	23
	Configuration Limitations	24
	Configuring Generic Routing Encapsulation Tunneling	24
	Configuring a GRE Tunnel	25
	Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly	25
<b>Part 3</b>	<b>Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP)</b>	
<b>Chapter 3</b>	<b>Understanding LAGs and LACP</b>	<b>29</b>
	Understanding Aggregated Ethernet Interfaces and LACP	29
	Link Aggregation Group	30
	Link Aggregation Control Protocol (LACP)	31
	Configuring Aggregated Ethernet LACP	32
	Configuring Link Aggregation	33
	Creating an Aggregated Ethernet Interface	34
	Configuring the VLAN Name and VLAN ID Number	34
	Configuring Aggregated Ethernet LACP	35
	Verifying the Status of a LAG Interface	36
	Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets	36
	Verifying the LACP Setup	36
	Verifying That LACP Packets Are Being Exchanged	37
	Troubleshooting an Aggregated Ethernet Interface	37
<b>Part 4</b>	<b>Integrated Routing and Bridging</b>	
<b>Chapter 4</b>	<b>Using Integrated Routing and Bridging</b>	<b>41</b>
	Understanding Integrated Routing and Bridging	41
	Configuring VLANs	43
	Configuring IRB Interfaces	44
	Example: Configuring Routing Between VLANs on One Switch	45
<b>Part 5</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 5</b>	<b>Configuration Statements for Interfaces</b>	<b>53</b>
	address	54
	alarm (chassis)	57
	configured-flow-control	58

	description (Interfaces) . . . . .	59
	ether-options . . . . .	60
	eui-64 . . . . .	61
	family . . . . .	62
	flow-control . . . . .	66
	fpc . . . . .	67
	hold-time (Physical Interface) . . . . .	68
	inet (interfaces) . . . . .	69
	inet6 (interfaces) . . . . .	70
	interface-range . . . . .	71
	interfaces . . . . .	73
	link-mode . . . . .	80
	link-speed . . . . .	81
	loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet) . . . .	82
	member . . . . .	82
	member-range . . . . .	83
	mtu . . . . .	84
	pic . . . . .	86
	rx-buffers . . . . .	87
	source . . . . .	88
	targeted-broadcast . . . . .	89
	traceoptions (Individual Interfaces) . . . . .	90
	tx-buffers . . . . .	91
	unit . . . . .	93
<b>Chapter 6</b>	<b>Configuration Statements for GRE . . . . .</b>	<b>95</b>
	destination (Tunnels) . . . . .	95
	tunnel . . . . .	96
	tunnel-port . . . . .	96
<b>Chapter 7</b>	<b>Configuration Statements for LAGs and LACP . . . . .</b>	<b>97</b>
	aggregated-devices . . . . .	98
	aggregated-ether-options . . . . .	99
	chassis . . . . .	101
	802.3ad . . . . .	102
	device-count . . . . .	103
	ethernet . . . . .	103
	force-up . . . . .	104
	lACP (802.3ad) . . . . .	105
	lACP (Aggregated Ethernet) . . . . .	106
	periodic . . . . .	106
<b>Chapter 8</b>	<b>Configuration Statements for IRB Interfaces . . . . .</b>	<b>107</b>
	irb (Interfaces) . . . . .	108
	I3-interface (VLAN) . . . . .	111
	vLans (OCX Series) . . . . .	112
<b>Chapter 9</b>	<b>Interfaces Operational Commands . . . . .</b>	<b>113</b>
	monitor interface . . . . .	114
	show interfaces diagnostics optics . . . . .	123

	show interfaces ge . . . . .	137
	show interfaces (GRE) . . . . .	152
	show interfaces queue . . . . .	159
	show interfaces xe . . . . .	199
<b>Chapter 10</b>	<b>LAGs and LACP Operational Commands . . . . .</b>	<b>217</b>
	show lacp interfaces . . . . .	218
	show lacp statistics interfaces (View) . . . . .	223

# List of Figures

Part 4	Integrated Routing and Bridging	
Chapter 4	Using Integrated Routing and Bridging .....	41
	Figure 1: IRB with One Switch .....	46





# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xi</b>
	Table 1: Notice Icons . . . . .	xiii
	Table 2: Text and Syntax Conventions . . . . .	xiii
<b>Part 1</b>	<b>Interfaces</b>	
<b>Chapter 1</b>	<b>Understanding Interfaces</b> . . . . .	<b>3</b>
	Table 3: Network Interface Types and Purposes . . . . .	3
	Table 4: Special Interface Types and Purposes . . . . .	4
	Table 5: LPM Table Combinations Using l2-profile-one with Junos OS 1Release 3.2X51-D10 . . . . .	16
	Table 6: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10 . . . . .	17
<b>Part 2</b>	<b>Generic Routing Encapsulation (GRE)</b>	
<b>Chapter 2</b>	<b>Understanding GRE</b> . . . . .	<b>21</b>
	Table 7: Firewall Filter Application Points for Tunneled Packets . . . . .	23
	Table 8: Features Not Supported with GRE . . . . .	24
<b>Part 4</b>	<b>Integrated Routing and Bridging</b>	
<b>Chapter 4</b>	<b>Using Integrated Routing and Bridging</b> . . . . .	<b>41</b>
	Table 9: Sample IRB Values . . . . .	42
	Table 10: Number of Supported IRBs/RVIs by Platform . . . . .	42
	Table 11: Components of the Multiple VLAN Topology . . . . .	46
<b>Part 5</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 9</b>	<b>Interfaces Operational Commands</b> . . . . .	<b>113</b>
	Table 12: Output Control Keys for the monitor interface Command . . . . .	114
	Table 13: Output Control Keys for the monitor interface traffic Command . . . . .	115
	Table 14: monitor interface Output Fields . . . . .	116
	Table 15: show interfaces diagnostics optics Output Fields . . . . .	123
	Table 16: show interfaces ge Output Fields . . . . .	138
	Table 17: GRE show interfaces Output Fields . . . . .	153
	Table 18: Layer 2 Overhead and Transmitted Packets or Byte Counts . . . . .	160
	Table 19: show interfaces queue Output Fields . . . . .	163
	Table 20: Byte Count by Interface Hardware . . . . .	166
	Table 21: show interfaces xe Output Fields . . . . .	200
<b>Chapter 10</b>	<b>LAGs and LACP Operational Commands</b> . . . . .	<b>217</b>

Table 22: show lacp interfaces Output Fields . . . . .	219
Table 23: show lacp statistics interfaces Output Fields . . . . .	223

# 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<sup>®</sup> 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:

- OCX1100

## Using the Examples in This Manual

---

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

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

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

## Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

## Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

## Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

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

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

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

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

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

## Documentation Feedback

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

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

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

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

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

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

## Opening a Case with JTAC

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

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

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



## PART 1

# Interfaces

- [Understanding Interfaces on page 3](#)



## CHAPTER 1

# Understanding Interfaces

- [Interfaces Overview on page 3](#)
- [Understanding Interface Naming Conventions on page 4](#)
- [Understanding Interface Ranges on page 5](#)
- [Understanding Management Interfaces on page 6](#)
- [Configuring the Interface Address on page 8](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13](#)
- [Configuring Ethernet Loopback Capability on page 14](#)
- [Configuring an LPM Table With Junos OS Release 13.2X51-D10 on page 15](#)
- [Monitoring Interface Status and Traffic on page 17](#)
- [Troubleshooting Network Interfaces on page 17](#)

## Interfaces Overview

---

Juniper Networks devices have two types of interfaces: network interfaces and special interfaces. This topic provides brief information about these interfaces. For additional information, see the *Junos OS Network Interfaces Library for Routing Devices*.

- [Network Interfaces on page 3](#)
- [Special Interfaces on page 4](#)

## Network Interfaces

Network interfaces connect to the network and carry network traffic. [Table 3 on page 3](#) lists the types of network interfaces supported.

**Table 3: Network Interface Types and Purposes**

Type	Purpose
Aggregated Ethernet interfaces	You can group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.
Ethernet Interfaces	You can configure Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet interfaces to connect to other servers, storage, and switches.

## Special Interfaces

Table 4 on page 4 lists the types of special interfaces supported.

**Table 4: Special Interface Types and Purposes**

Type	Purpose
Console port	Each device has a serial console port, labeled <b>CON</b> or <b>CONSOLE</b> , for connecting tty-type terminals to the switch. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch.
Loopback interface	A software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	<p>The management Ethernet interface provides an out-of-band method for connecting to a standalone switch and QFabric system.</p> <p><b>NOTE:</b> On OCX Series switches, the em0 management interface always has the status <b>up</b> in <b>show</b> command outputs, even if the physical port is empty. The me0 interface is a virtual interface between Junos and the host operating system, therefore its status is independent from the status of the physical port.</p>

- Related Documentation**
- [Understanding Interface Naming Conventions on page 4](#)
  - [Understanding Management Interfaces on page 6](#)

## Understanding Interface Naming Conventions

The OCX Series device uses a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos OS. This topic provides brief information about the naming conventions used for interfaces on the OCX Series.

This topic describes:

- [Physical Part of an Interface Name on page 4](#)
- [Wildcard Characters in Interface Names on page 5](#)

### Physical Part of an Interface Name

Interfaces in Junos OS are specified as follows:

*type-fpc/pic/port*

The convention is as follows:

- *type*—The OCX Series device interfaces use the following media types:
  - **xe**—10-Gigabit Ethernet interface
  - **et**—40-Gigabit Ethernet interface

- **em**—Management interface
- **fpc**—Flexible PIC Concentrator. OCX Series interfaces use the following convention for the FPC number in interface names:

- On standalone OCX Series switches, the FPC number is always 0.

The FPC number indicates the slot number of the line card that contains the physical interface.

- **pic**—The OCX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
  - PIC 0 provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces.
- **port**—Interfaces use the following convention for port numbers:
  - PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 1 through 48 and 6 40-Gbps QSFP+ ports labeled 49 through 54.

## Wildcard Characters in Interface Names

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the **interface-name** option to specify groups of interface names without having to type each name individually. You must enclose all wildcard characters except the asterisk (\*) in quotation marks (" ").

### Related Documentation

- [Interfaces Overview on page 3](#)
- [Understanding Management Interfaces on page 6](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## Understanding Interface Ranges

You can use the interface ranges to group interfaces of the same type that share a common configuration profile. This helps reduce the time and effort in configuring interfaces. The configurations common to all the interfaces can be included in the interface range definition.

The interface range definition contains the name of the interface range defined, the names of the individual member interfaces that do not fall in a series of interfaces, a range of interfaces defined in the member range, and the configuration statements common to all the interfaces. An interface range defined with member ranges and individual members but without any common configurations is also a valid definition.



**NOTE:** The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces. OCX Series switches do not support Fibre Channel interfaces.

The common configurations defined in the interface range will be overridden by the local configuration.

The defined interface ranges can be used at places where the **interface** statement is used in the following configuration hierarchies:



**NOTE:** These statements are not supported on OCX Series switches:

- **protocols isis interface**
- **protocols sflow interfaces**

**Related  
Documentation**

- *Interfaces Overview*
- [Interfaces Overview on page 3](#)
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13](#)
- [Configuring Link Aggregation on page 33](#)
- *Configuring a Layer 3 Logical Interface*
- *Junos OS Network Interfaces Library for Routing Devices*
- [interface-range on page 71](#)

---

## Understanding Management Interfaces

You use management interfaces to access devices remotely. Typically, a management interface is not connected to the in-band network, but is connected to a device in the internal network. Through a management interface, you can access the device over the network using utilities such as **ssh** and **telnet** and configure it from anywhere, regardless of its physical location. As a security feature, users cannot log in as **root** through a management interface. To access the device as **root**, you must use the console port. You can also use **root** to log in using SSH.



**NOTE:** Before you can use management interfaces, you must configure the logical interfaces with valid IP addresses. Juniper Networks does not support configuring two management interfaces in the same subnet.

Management interface port ranges vary based on device type:

- QFX3500 devices:

The valid port range for a management interface (**me**) on a QFX3500 device is between 0 and 6, with a total of seven available ports. On a QFX3500 standalone switch, however, you can only configure **me0** and **me1** as management interfaces. The management interfaces are labeled **C0** and **C1**, and they correspond to **me0** and **me1**. On a QFX3500 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**.

- QFX3600 devices:

There are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0S** and **C1S**). On a QFX3600 standalone switch, the RJ-45 management interfaces and SFP management interfaces correspond to **me0** and **me1**. On a QFX3600 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**. Each pair of management interfaces correspond to one Ethernet interface—for example, both RJ-45 management interfaces (labeled **C0** and **C0S**) can correspond to **me0**, and both SFP management interfaces (labeled **C1** and **C1S**) can correspond to **me1**. By default, both RJ-45 management interfaces are active. If you insert an SFP interface into the SFP management port (**C0S**, for example), the SFP interface would become the active management interface, and the corresponding RJ-45 management interface (**C0**) is disabled.



**NOTE:** On a QFX3600 device, you can use either the RJ-45 or the SFP management interfaces, but not both at the same time.

- On QFX5100 and EX4600 switches, there is one RJ-45 management interface (labeled **C0**) and one SFP management interface (labeled **C1**), and they correspond to **em0** and **em1**. You can use both management interfaces simultaneously.
- On OCX Series switches:

There is one RJ-45 management interface (labeled **MGMT**), which corresponds to **em0**. The **em0** interface always has the status **up** in show command outputs, even if the physical port is empty. The **me0** interface is a virtual interface between Junos and the host operating system, therefore its status is independent from the status of the physical port.

- QFabric system:

On a QFabric system, there are management interfaces on the Node devices, Interconnect devices, and Director devices. However, you cannot access the management interfaces on the Node devices or Interconnect devices directly. You can only manage and configure these devices using the Director device. You can connect to the management interface over the network using utilities such as SSH.

For information on how to use management interfaces on a QFabric system, see *Performing the QFabric System Initial Setup on a QFX3100 Director Group* and *Gaining Access to the QFabric System Through the Default Partition*.

- Related Documentation
- [Interfaces Overview](#)
  - [Interfaces Overview on page 3](#)

## Configuring the Interface Address

---

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vpls** families, you never configure an address.



**NOTE:** The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, include the **address** statement:

```
address address {  
    broadcast address;  
    destination address;  
    destination-profile name;  
    eui-64;  
    preferred;  
    primary;  
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **address** statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

- Broadcast address for the interface subnet—Specify this in the **broadcast** statement; this applies only to Ethernet interfaces, such as the management interface **fxp0**, **em0**, or **me0** the Fast Ethernet interface, and the Gigabit Ethernet interface.
- Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the **destination** statement.
- PPP properties to the remote end—Specify this in the **destination-profile** statement. You define the profile at the [edit access group-profile *name* **ppp**] hierarchy level (for point-to-point interfaces only).
- Whether the router or switch automatically generates the host number portion of interface addresses—The **eui-64** statement applies only to interfaces that carry IPv6 traffic, in which the prefix length of the address is 64 bits or less, and the low-order 64



bits of the address are zero. This option does not apply to the loopback interface (**lo0**) because IPv6 addresses configured on the loopback interface must have a 128-bit prefix length.

- Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet.

By default, the preferred address is the lowest-numbered address on the subnet. To override the default and explicitly configure the preferred address, include the **preferred** statement when configuring the address.

- Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you send packets from an interface where the destination provides no information about the subnet (for example, some **ping** commands).

By default, the primary address on an interface is the lowest-numbered non-127 (in other words, non-loopback) preferred address on the interface. To override the default and explicitly configure the preferred address, include the **primary** statement when configuring the address.



**NOTE:** If you configure a duplicate IP address on an interface, even when the earlier interface with that IP address is disabled, a Warning message is added to the syslog and not displayed on the screen. Do not configure the same IP address of a disabled interface on another interface.

- [Configuring Interface IPv4 Addresses on page 9](#)
- [Configuring Interface IPv6 Addresses on page 12](#)

## Configuring Interface IPv4 Addresses

You can configure router or switch interfaces with a 32-bit IP version 4 (IPv4) address and optionally with a destination prefix, sometimes called a *subnet mask*. An IPv4 address utilizes a 4-octet dotted decimal address syntax (for example, **192.16.1.1**). An IPv4 address with destination prefix utilizes a 4-octet dotted decimal address syntax with a destination prefix appended (for example, **192.16.1.1/30**).

To configure an IPv4 address on routers and switches running Junos OS, use the **edit interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.



**NOTE:** Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.

## Operational Behavior of Interfaces When the Same IPv4 Address Is Assigned to Them

You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point.



**NOTE:** By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.



**NOTE:** If you configure the same IP address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration. The remaining IP address configurations are ignored, leaving some interfaces without an assigned address. Interfaces without an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

In the following example, the IP address configuration for interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
}
```

The following examples show the sample configuration of assigning the same IPv4 address to implicitly and explicitly point-to-point interfaces, and their corresponding **show interfaces terse** command outputs to see their operational status.

### Configuring same IPv4 address on implicitly PPP interfaces:

```
[edit]
user@host# show
ge-0/1/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
```

```

    }
  }
}
ge-3/0/1 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
}

```

The sample output shown below for the above configuration reveals that only **ge-0/1/0.0** was assigned the same IPv4 address **200.1.1.1/24** and its **link** state was **up**, while **ge-3/0/1.0** was not assigned the IPv4 address, though its **link** state was **up**, which means that it will be operational only when it gets a unique IPv4 address other than **200.1.1.1/24**.

```

user@host> show interfaces terse ge*
Interface           Admin Link Proto  Local          Remote
      ge-0/1/0             up   up
      ge-0/1/0.0           up   up   inet    200.1.1.1/24
                        multiservice
      ge-0/1/1             up   down
      ge-3/0/0             up   down
      ge-3/0/1             up   up
      ge-3/0/1.0           up   up   inet
                        multiservice

```

Configuring same IPv4 address on explicitly PPP interfaces:

```

[edit]
user@host# show
so-0/0/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
so-0/0/3 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}

```

The sample output shown below for the above configuration reveals that both **so-0/0/0.0** and **so-0/0/3.0** were assigned the same IPv4 address **200.1.1.1/24** and that their **link** states were **down**, which means that to make them operational at least one of them will have to be configured with a unique IPv4 address other than **200.1.1.1/24**.

```

user@host> show interfaces terse so*
Interface           Admin Link Proto  Local          Remote
so-0/0/0             up   up
so-0/0/0.0           up   down inet    200.1.1.1/24
so-0/0/1             up   up
so-0/0/2             up   down

```

so-0/0/3	up	up	
so-0/0/3.0	up	down	inet 200.1.1.1/24
so-1/1/0	up	down	
so-1/1/1	up	down	
so-1/1/2	up	up	
so-1/1/3	up	up	
so-2/0/0	up	up	
so-2/0/1	up	up	
so-2/0/2	up	up	
so-2/0/3	up	down	

## Configuring Interface IPv6 Addresses



**NOTE:** IPv6 is not currently supported for the QFX Series.

You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values.

You assign a 128-bit IPv6 address to an interface by including the **address** statement:

```
address aaaa:bbbb:...:zzzz/nn;
```



**NOTE:** You cannot configure a subnet zero IPv6 address because RFC 2461 reserves the subnet-zero address for anycast addresses, and Junos OS complies with the RFC.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6]

The double colon (::) represents all bits set to 0, as shown in the following example:

```
interfaces fe-0/0/1 {
  unit 0 {
    family inet6 {
      address fec0:1:1::2/64;
    }
  }
}
```



**NOTE:** You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

### Related Documentation

- *Configuring IPCP Options*
- *Configuring Default, Primary, and Preferred Addresses and Interfaces*

## Configuring Gigabit and 10-Gigabit Ethernet Interfaces

Devices include a factory default configuration that:

- Enables all 10-Gigabit Ethernet network interfaces on the switch
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching**
- Configures Storm Control on all 10-Gigabit Ethernet network interfaces

This topic describes:

- [Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces on page 13](#)
- [Configuring the IP Options on page 13](#)

### Configuring the Link Settings for Gigabit Ethernet and 10-Gigabit Ethernet Interfaces

Devices include a factory default configuration that enables 10-Gigabit Ethernet and interfaces with applicable link settings.

The following default configurations are available on 10-Gigabit Ethernet interfaces:

- The speed for 10-Gigabit Ethernet interfaces is set to 10 Gbps by default. The speed cannot be configured.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is not supported.

The **ether-options** statement enables you to modify the following options:

- **802.3ad**—Specify an aggregated Ethernet bundle for 10-Gigabit Ethernet interfaces.
- **loopback**—Enable or disable a loopback interface for 10-Gigabit Ethernet interfaces.

To set **ether-options** for 10-Gigabit Ethernet interfaces:

```
[edit]
user@switch# set interfaces interface-name ether-options
```

### Configuring the IP Options

To specify an IP address for the logical unit:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

#### Related Documentation

- [Monitoring Interface Status and Traffic on page 17](#)
- [show interfaces xe on page 199](#)
- *show interfaces ge-*
- *speed*

- [Understanding Interface Naming Conventions on page 4](#)
- [Interfaces Overview on page 3](#)

## Configuring Ethernet Loopback Capability

---

To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

To return to the default—that is, to disable loopback mode—delete the **loopback** statement from the configuration:

```
[edit]  
user@switch# delete interfaces interface-name ether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

```
no-loopback;
```

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- **[edit interfaces *interface-name* aggregated-ether-options]**
- **[edit interfaces *interface-name* ether-options]**

### Related Documentation

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13](#)

## Configuring an LPM Table With Junos OS Release 13.2X51-D10

---

In addition to choosing a profile, you can further optimize memory allocation for LPM table entries by configuring how many IPv6 addresses with prefixes in the range /65 through /127 you want to store. If you want to use more than 16 IPv6 addresses with prefixes in this range, you must enter and commit the following statement:

[edit]

```
user@switch# set chassis forwarding-options profile-name num-65-127-prefix value
```

in which *value* can be a value in the range 1 through 128. Each increment adds support for 16 IPv6 addresses with prefixes between /65 and /127, for a maximum of 2048 such addresses (16 x 128 = 2048). The system supports 16 of these addresses by default, so to increase the number of supported addresses, you must enter a value of 2 or greater. For example, if you enter **2**, the system will support 32 IPv6 addresses with prefixes in the range /65 through /127.



**NOTE:** When you configure the `num-65-127-prefix` value, all the data interfaces on the switch restart. The management interfaces are unaffected.

The LPM table is shared, and each increment that you add for IPv6 addresses with prefixes in the range /65 through /127 reduces the number of forwarding table entries that are available for IPv4 addresses and IPv6 addresses with prefixes less than /65.

[Table 5 on page 16](#) provides examples of valid combinations that the LPM table can store, also using the **l2-profile-one** profile. Once again, each row in the table represents a case in which the table is full and cannot accommodate any more entries.

**Table 5: LPM Table Combinations Using l2-profile-one with Junos OS 1Release 3.2X51-D10**

IPv4 entries	IPv6 Entries (prefix <= 64)	IPv6 Entries (prefix >= 65)	num-65-127-prefix
16K	0K	16	1 (default)
0K	8K	16	1 (default)
8K	4K	16	1 (default)
4K	4K	1K	<b>64</b>
2K	5K	1K	<b>64</b>
0K	6K	1K	<b>64</b>
4K	2K	2K	<b>128</b>
2K	3K	2K	<b>128</b>
0K	4K	2K	<b>128</b>

[Table 6 on page 17](#) provides examples of valid combinations that the LPM table can store when you use the **lpm-profile** profile. As before, each row represents a case in which the table is full and cannot accommodate any more entries.



Table 6: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10

IPv4 entries	IPv6 Entries (prefix <= 64)	IPv6 Entries (prefix >= 65)	num-65-127-prefix
128K	0K	16	1 (default)
0K	8K	16	1 (default)
8K	4K	16	1 (default)
4K	4K	1K	64
2K	5K	1K	64
0K	6K	1K	64
4K	2K	2K	128
2K	3K	2K	128
0K	4K	2K	128

**Related Documentation**

- [Understanding the Unified Forwarding Table](#)
- [Configuring the Unified Forwarding Table](#)

## Monitoring Interface Status and Traffic

**Purpose** View interface status to monitor interface bandwidth utilization and traffic statistics.

- Action**
- To view interface status for all the interfaces, enter [show interfaces xe](#).
  - To view status and statistics for a specific interface, enter [show interfaces xe interface-name](#).
  - To view status and traffic statistics for all interfaces, enter either [show interfaces xe detail](#) or [show interfaces xe extensive](#).

**Meaning** For details about output from the CLI commands, see [show interfaces xe](#).

## Troubleshooting Network Interfaces

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down

**Problem Description:** The switch has an SFP or SFP+ module installed. The interface on the port in which an SFP or SFP+ transceiver is installed is down.

**Symptoms:** When you check the status with the CLI command [show interfaces interface-name](#), the disabled port is not listed.

**Cause** By default, the SFP or SFP+ module operates in the 10-Gigabit Ethernet mode and supports only SFP or SFP+ transceivers. The operating mode for the module is incorrectly set.

**Solution** Only SFP or SFP+ transceivers can be installed in SFP or SFP+ modules. You must configure the operating mode of the SFP or SFP+ module to match the type of transceiver you want to use. For SFP+ transceivers, configure 10-Gigabit Ethernet operating mode.

## PART 2

# Generic Routing Encapsulation (GRE)

- [Understanding GRE on page 21](#)



## CHAPTER 2

# Understanding GRE

- [Understanding Generic Routing Encapsulation on page 21](#)
- [Configuring Generic Routing Encapsulation Tunneling on page 24](#)
- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 25](#)

## Understanding Generic Routing Encapsulation

---

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets.

This topic describes:

- [Overview of GRE on page 21](#)
- [GRE Tunneling on page 22](#)
- [Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches on page 23](#)
- [Configuration Limitations on page 24](#)

### Overview of GRE

GRE encapsulates data packets and redirects them to a device that de-encapsulates them and routes them to their final destination. This allows the source and destination switches to operate as if they have a virtual point-to-point connection with each other (because the outer header applied by GRE is transparent to the encapsulated payload packet). For example, GRE tunnels allow routing protocols such as RIP and OSPF to forward data packets from one switch to another switch across the Internet. In addition, GRE tunnels can encapsulate multicast data streams for transmission over the Internet.

GRE is described in RFC 2784 (obsoletes earlier RFCs 1701 and 1702). The switches support RFC 2784, but not completely. (For a list of limitations, see [“Configuration Limitations” on page 24.](#))

As a *tunnel source router*, the switch encapsulates a payload packet for transport through the tunnel to a destination network. The payload packet is first encapsulated in a GRE packet, and then the GRE packet is encapsulated in a delivery protocol. The switch performing the role of a *tunnel remote router* extracts the tunneled packet and forwards the packet to its destination. Note that you can use one firewall term to terminate many GRE tunnels on a QFX5100 switch.

## GRE Tunneling

Data is routed by the system to the GRE endpoint over routes established in the route table. (These routes can be statically configured or dynamically learned by routing protocols such as RIP or OSPF.) When a data packet is received by the GRE endpoint, it is de-encapsulated and routed again to its destination address.

GRE tunnels are *stateless*—that is, the endpoint of the tunnel contains no information about the state or availability of the remote tunnel endpoint. Therefore, the switch operating as a tunnel source router cannot change the state of the GRE tunnel interface to down if the remote endpoint is unreachable.

For details about GRE tunneling, see:

- [Encapsulation and De-Encapsulation on the Switch on page 22](#)
- [Number of Source and Destination Tunnels Allowed on a Switch on page 22](#)
- [Class of Service on GRE Tunnels on page 23](#)
- [Applying Firewall Filters to GRE Traffic on page 23](#)

---

### Encapsulation and De-Encapsulation on the Switch

Encapsulation—A switch operating as a tunnel source router encapsulates and forwards GRE packets as follows:

1. When a switch receives a data packet (payload) to be tunneled, it sends the packet to the tunnel interface.
2. The tunnel interface encapsulates the data in a GRE packet and adds an outer IP header.
3. The IP packet is forwarded on the basis of the destination address in the outer IP header.

De-encapsulation—A switch operating as a tunnel remote router handles GRE packets as follows:

1. When the destination switch receives the IP packet from the tunnel interface, the outer IP header and GRE header are removed.
2. The packet is routed based on the inner IP header.

---

### Number of Source and Destination Tunnels Allowed on a Switch

QFX5100 and OCX Series switches support as many as 512 GRE tunnels, including tunnels created with a firewall filter. That is, you can create a total of 512 GRE tunnels, regardless of which method you use.

EX switches support as many as 500 GRE tunnels between switches transmitting IPv4 or IPv6 payload packets over GRE. If a passenger protocol in addition to IPv4 and IPv6 is used, you can configure up to 333 GRE tunnels between the switches.

An EX switch can have a maximum of 20 tunnel source IP addresses configured, and each tunnel source IP can be configured with up to 20 destination IP addresses on a second switch. As a result, the two connected switches can have a maximum of 400 GRE tunnels. If the first switch is also connected to a third switch, the possible maximum number of tunnels is 500.

### Class of Service on GRE Tunnels

When a network experiences congestion and delay, some packets might be dropped. Junos OS class of service (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs and thereby set rules for packet loss. For details about CoS, see [Junos OS CoS for EX Series Switches Overview](#).

The following CoS components are available on a switch operating as a GRE tunnel source router or GRE tunnel remote router:

- At the GRE tunnel source—On a switch operating as a tunnel source router, you can apply CoS classifiers on an *ingress port* or on a *GRE port*, with the following results on CoS component support on tunneled packets:
  - Schedulers only—Based on the CoS classification on the ingress port, you can apply CoS schedulers on a GRE port of the switch to define output queues and control the transmission of packets through the tunnel after GRE encapsulation. However, you cannot apply CoS rewrite rules to these packets.
  - Schedulers and rewrite rules—Depending on the CoS classification on the GRE port, you can apply both schedulers and rewrite rules to the encapsulated packets transmitted through the tunnel.
- At the GRE tunnel endpoint—When the switch is a tunnel remote router, you can apply CoS classifiers on the GRE port and schedulers and rewrite rules on the egress port to control the transmission of a de-encapsulated GRE packet out from the egress port.

### Applying Firewall Filters to GRE Traffic

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a switch. (For details, see [Firewall Filters for EX Series Switches Overview](#).) Because of the encapsulation and de-encapsulation performed by GRE, you are constrained as to where you can apply a firewall filter to filter tunneled packets and which header will be affected. [Table 7 on page 23](#) identifies these constraints.

**Table 7: Firewall Filter Application Points for Tunneled Packets**

Endpoint Type	Ingress Interface	Egress Interface
Source (encapsulating)	inner header	outer header
Remote (de-encapsulating)	Cannot filter packets on ingress interface	inner header

### Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches

You can also use a firewall filter to de-encapsulate GRE traffic on switches. This feature provides significant benefits in terms of scalability, performance, and flexibility because

you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source IP addresses with one firewall term. See *Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch* for information about how to configure a firewall filter for this purpose.

## Configuration Limitations

Table 8 on page 24 lists features that are not supported with GRE.

**Table 8: Features Not Supported with GRE**

EX Switches	QFX Switches
MPLS over GRE tunnels	MPLS over GRE tunnels
GRE keepalives	GRE keepalives
GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets	GRE keys, payload packet fragmentation, and sequence numbers for fragmented packets
BGP dynamic tunnels	BGP dynamic tunnels
Outer IP address must be IPv4	Outer IP address must be IPv4
Virtual routing instances	
Bidirectional Forwarding Detection (BFD) protocol over GRE distributed mode	
OSPF limitation—Enabling OSPF on a GRE interface creates two equal-cost routes to the destination: one through the Ethernet network or uplink interface and the other through the tunnel interface. If data is routed through the tunnel interface, the tunnel might fail. To keep the interface operational, we recommend that you use a static route, disable OSPF on the tunnel interface, or configure the peer not to advertise the tunnel destination over the tunnel interface.	

- Related Documentation**
- *Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)*
  - [Configuring Generic Routing Encapsulation Tunneling on page 24](#)
  - *Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch*

## Configuring Generic Routing Encapsulation Tunneling

Generic routing encapsulation (GRE) provides a private, secure path for transporting packets through an otherwise public network by encapsulating (or tunneling) the packets. GRE tunneling is accomplished through tunnel endpoints that encapsulate or de-encapsulate traffic.

You can also use a firewall filter to de-encapsulate GRE traffic on QFX5100 and OCX Series switches. This feature provides significant benefits in terms of scalability, performance, and flexibility because you don't need to create a tunnel interface to perform the de-encapsulation. For example, you can terminate many tunnels from multiple source



IP addresses with one firewall term. For more information on this feature, see *Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch*.

This topic describes:

1. [Configuring a GRE Tunnel on page 25](#)

## Configuring a GRE Tunnel

To configure a GRE tunnel interface:

1. Create a GRE interface with a unit number and address:

```
[edit interfaces]
user@switch# set gr-0/0/0 unit number family inet address
```



**NOTE:** The base name of the interface must be `gr-0/0/0`.

This is a pseudo interface, and the address you specify can be any IP address. The routing table must specify `gr-0/0/0.x` as the outgoing interface for any packets that will be tunneled.

If you configure a GRE interface on a QFX5100 switch that is a member of a Virtual Chassis and later change the Virtual Chassis member number of the switch, the name of the GRE interface does not change in any way (because it is a pseudo interface). For example, if you change the member number from `0` to `5`, the GRE interface name does *not* change from `gr-0/0/0.x` to `gr-5/0/0.x`.

2. Specify the tunnel source address for the logical interface:

```
[edit interfaces]
user@switch# set gr-0/0/0 unit number tunnel source source-address
```

3. Specify the destination address:

```
[edit interfaces]
user@switch# set gr-0/0/0 unit number tunnel destination destination-address
```

The destination address must be reachable through static or dynamic routing. If you use static routing, you must get the destination MAC address (for example, by using `ping`) before user traffic can be forwarded through the tunnel.

### Related Documentation

- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 25](#)
- [Understanding Generic Routing Encapsulation on page 21](#)
- [Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch](#)

## Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

**Purpose** Verify that the generic routing encapsulation (GRE) interface is sending tunneled traffic.

**Action** Display status information about the specified GRE interface by using the command [show interfaces](#).

```
user@switch> show interfaces gr-0/0/0.0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
  Flags: None
  Addresses, Flags: Is-Primary
    Local: 1.10.1.1
```

**Meaning** The output indicates that the GRE interface gr-0/0/0 is up. The output displays the name of the physical interface and the traffic statistics for this interface---the number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.

**Related Documentation**

- *Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)*

## PART 3

# Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP)

- [Understanding LAGs and LACP on page 29](#)



## CHAPTER 3

# Understanding LAGs and LACP

- [Understanding Aggregated Ethernet Interfaces and LACP on page 29](#)
- [Configuring Aggregated Ethernet LACP on page 32](#)
- [Configuring Link Aggregation on page 33](#)
- [Verifying the Status of a LAG Interface on page 36](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 36](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 37](#)

### Understanding Aggregated Ethernet Interfaces and LACP

---

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single, aggregated Ethernet interface, also known as a *link aggregation group (LAG)* or *bundle*.

Link aggregation is used to aggregate Ethernet interfaces between two devices. You can create a LAG between a Juniper Networks device and a router, switch, aggregation switch, server, or other devices. The aggregated Ethernet interfaces that participate in a LAG are called member links. Because a LAG is composed of multiple member links, even if one member link fails, the LAG continues to carry traffic over the remaining links.



**NOTE:** On QFX5100 and EX4600 standalone switches and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.



**NOTE:** The QFX5200 switches do not support mixed rate aggregated Ethernet bundles.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard and is used as a discovery protocol.



**NOTE:** To ensure load balancing across the aggregated Ethernet (AE) interfaces on a redundant server Node group, the members of the AE must be equally distributed across the redundant server Node group.



**NOTE:** During a network Node group switchover, traffic might be dropped for a few seconds.

- [Link Aggregation Group on page 30](#)
- [Link Aggregation Control Protocol \(LACP\) on page 31](#)

## Link Aggregation Group

To create a LAG:

1. Create a logical aggregated Ethernet interface.
2. Define the parameters associated with the logical aggregated Ethernet interface, such as a logical unit, interface properties, and Link Aggregation Control Protocol (LACP).
3. Define the member links to be contained within the aggregated Ethernet interface—for example, two 10-Gigabit Ethernet interfaces.
4. Configure LACP for link detection.

Keep in mind these hardware and software guidelines:

- Up to 32 Ethernet interfaces can be grouped to form a LAG on a redundant server Node group, a server Node group, and a network Node group on a QFabric system. Up to 48 LAGs are supported on redundant server Node groups and server Node groups on a QFabric system, and up to 128 LAGs are supported on network Node groups on a QFabric system. You can configure LAGs across Node devices in redundant server Node groups, server Node groups, and network Node groups.



**NOTE:** If you try to commit a configuration containing more than 32 Ethernet interfaces in a LAG, you will receive an error message saying that the group limit of 32 has been exceeded, and the configuration checkout has failed.

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and up to 448 LAGs are supported on QFX3500, QFX3600, EX4600, and OCX Series switches, and up to 1,000 LAGs are supported on QFX5100 switches.



**NOTE:** If you try to commit a configuration containing more than 64 Ethernet interfaces in a LAG, you will receive an error message saying that the group limit of 64 has been exceeded, and the configuration checkout has failed.

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and up to 144 LAGs are supported on QFX10002-36Q switches, and up to 288 LAGs are supported on QFX10002-72Q switches.
- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed and be in full-duplex mode.



**NOTE:** On a QFX5100, EX4600, QFX10002 standalone switch or QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure mixed rate aggregated Ethernet bundles (LAGs with different link speeds). OCX Series switches do not support LAGs with different speeds.



**NOTE:** Junos OS assigns a unique ID and port priority to each port. The ID and priority are not configurable.

- QFabric systems support a special LAG called an FCoE LAG, which enables you to transport FCoE traffic and regular Ethernet traffic (traffic that is not FCoE traffic) across the same link aggregation bundle. Standard LAGs use a hashing algorithm to determine which physical link in the LAG is used for a transmission, so communication between two devices might use different physical links in the LAG for different transmissions. An FCoE LAG ensures that FCoE traffic uses the same physical link in the LAG for requests and replies in order to preserve the virtual point-to-point link between the FCoE device converged network adapter (CNA) and the FC SAN switch across a QFabric system Node device. An FCoE LAG does not provide load balancing or link redundancy for FCoE traffic. However, regular Ethernet traffic uses the standard hashing algorithm and receives the usual LAG benefits of load balancing and link redundancy in an FCoE LAG. See *Understanding FCoE LAGs* for more information.

## Link Aggregation Control Protocol (LACP)

LACP is one method of bundling several physical interfaces to form one logical aggregated Ethernet interface. The LACP mode can be active or passive. The transmitting link is known as the *actor*, and the receiving link is known as the *partner*. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode. You can configure Ethernet links to actively transmit protocol data units (PDUs), or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. You can configure both VLAN-tagged and untagged aggregated Ethernet interfaces without LACP enabled. LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the LAG without user intervention.

- Link monitoring to check whether both ends of the bundle are connected to the correct group.

When a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server may not be able to exchange LACP PDUs. In such a situation you can configure an interface to be in the **up** state even if no PDUs are exchanged. Use the **force-up** statement to configure an interface when the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When there are no received PDUs, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

#### Related Documentation

- [Configuring Link Aggregation on page 33](#)
- *Configuring an FCoE LAG*
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- *Example: Configuring an FCoE LAG on a Redundant Server Node Group*
- [Verifying the Status of a LAG Interface on page 36](#)
- *Junos OS Network Interfaces Library for Routing Devices*

---

## Configuring Aggregated Ethernet LACP

---

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

---



To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp active
```

2. Specify the interval and speed at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic interval
```

For example, to specify the interval as fast, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic fast
```

## Configuring Link Aggregation

Use the link aggregation feature to aggregate one or more links to form a virtual link or aggregation group. The MAC client can treat this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases link availability.



**NOTE:** An interface with an already configured IP address cannot form part of the aggregation group.



**NOTE:** On QFX5100 and EX4600 standalone switches and on QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.

1. [Creating an Aggregated Ethernet Interface on page 34](#)
2. [Configuring the VLAN Name and VLAN ID Number on page 34](#)
3. [Configuring Aggregated Ethernet LACP on page 35](#)

## Creating an Aggregated Ethernet Interface

To create an aggregated Ethernet interface:

1. Specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count device-count
```

For example, to specify 5:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count
```

2. Specify the minimum number of links for the aggregated Ethernet interface (aex), that is, the defined bundle, to be labeled “up”:



**NOTE:** By default only one link must be up for the bundle to be labeled “up”.

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links minimum-links
```

For example, to specify 5:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options minimum-links 5
```

3. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed link-speed
```

For example, to specify 10g:

```
[edit interfaces]
user@switch# set interface-name aggregated-ether-options link-speed 10g
```

4. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set interface-name ether-options 802.3ad aex
user@switch# set interface-name ether-options 802.3ad aex
```

## Configuring the VLAN Name and VLAN ID Number



**NOTE:** VLANs are not supported on OCX Series switches.

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

For example, 100.

## Configuring Aggregated Ethernet LACP

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs).

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.



**NOTE:** Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable the LACP mode:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp mode
```

For example, to specify the mode as active, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp active
```

2. Specify the interval and speed at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic interval
```

For example, to specify the interval as fast, execute the following command:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic fast
```

### Related Documentation

- [Understanding Interface Naming Conventions](#)
- [Understanding Interface Naming Conventions on page 4](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch](#)
- [Verifying the Status of a LAG Interface on page 36](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 36](#)
- [show lacp statistics interfaces \(View\) on page 223](#)

## Verifying the Status of a LAG Interface

**Purpose** Verify that a link aggregation group (LAG) (**ae0**) has been created on the switch.

**Action** To verify that the **ae0** LAG has been created:

```
[edit interfaces]
show interfaces ae0 terse
```

Interface	Admin	Link	Proto	Local	Remote
ae0	up	up			
ae0.0	up	up	inet	10.10.10.2/24	

**Meaning** The output confirms that the **ae0** link is up and shows the family and IP address assigned to this link.

### Related Documentation

- [Configuring Link Aggregation on page 33](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 36](#)
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*
- [show lacp statistics interfaces \(View\) on page 223](#)

## Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

Verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

1. [Verifying the LACP Setup on page 36](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 37](#)

### Verifying the LACP Setup

**Purpose** Verify that the LACP has been set up correctly.

**Action** To verify that LACP has been enabled as active on one end:

```
user@switch>show lacp interfaces xe-0/0/0
```

```
Aggregated interface: ae0
```

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
xe-0/1/0	Actor	No	Yes	No	No	No	Yes	Fast	Active
xe-0/1/0	Partner	No	Yes	No	No	No	Yes	Fast	Passive

LACP protocol:	Receive State	Transmit State	Mux	State
xe-0/1/0	Defaulted	Fast periodic		Detached

**Meaning** This example shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, one side must be set as active in order for the bundled link to be up.

## Verifying That LACP Packets Are Being Exchanged

**Purpose** Verify that LACP packets are being exchanged between interfaces.

**Action** Use the `show lacp statistics interfaces interface-name` command to display LACP BPDU exchange information.

`show lacp statistics interfaces ae0`

Aggregated interface: ae0

LACP Statistics:	LACP Rx	LACP Tx	Unknown Rx	Illegal Rx
xe-0/0/2	1352	2035	0	0
xe-0/0/3	1352	2056	0	0

**Meaning** The output here shows that the link is up and that PDUs are being exchanged.

- Related Documentation**
- [Configuring Link Aggregation on page 33](#)
  - [Verifying the Status of a LAG Interface on page 36](#)
  - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch](#)
  - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch](#)
  - [show lacp statistics interfaces \(View\) on page 223](#)

## Troubleshooting an Aggregated Ethernet Interface

**Problem** **Description:** The `show interfaces terse` command shows that the LAG is down.

**Solution** Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet-switching (Layer 2 LAG) or family inet (Layer 3 LAG).



**NOTE:** Layer 2 LAGs are not supported on OCX Series switches.

- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch.

**Related  
Documentation**

- [Verifying the Status of a LAG Interface on page 36](#)
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*

## PART 4

# Integrated Routing and Bridging

- [Using Integrated Routing and Bridging on page 41](#)





## CHAPTER 4

# Using Integrated Routing and Bridging

- [Understanding Integrated Routing and Bridging on page 41](#)
- [Configuring VLANs on page 43](#)
- [Configuring IRB Interfaces on page 44](#)
- [Example: Configuring Routing Between VLANs on One Switch on page 45](#)

### Understanding Integrated Routing and Bridging

---

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). VLANs limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN. For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs, you normally you need a router that connects the VLANs. However, you can accomplish this forwarding on a switch without using a router by configuring an integrated routing and bridging (IRB) interface. (These interfaces are also called routed VLAN interfaces, or RVIs). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

An IRB is a special type of Layer 3 virtual interface named **vlan**. Like normal Layer 3 interfaces, the **vlan** interface needs a logical unit number with an IP address. In fact, to be useful an IRB needs at least two logical units and two IP addresses—you must create units with addresses in each of the subnets associated with the VLANs between which you want traffic to be routed. That is, if you have two VLANs (for example, VLAN **red** and VLAN **blue**) with corresponding subnets, your IRB must have a logical unit with an address in the subnet for **red** and a logical unit with an address in the subnet for **blue**. The switch automatically creates direct routes to these subnets and uses these routes to forward traffic between VLANs.

[Table 9 on page 42](#) shows values you might use when configuring an IRB:

Table 9: Sample IRB Values

Property	Settings
VLAN names and tags (IDs)	<b>blue</b> , ID 100 <b>red</b> , ID 200
Subnets associated with VLANs	<b>blue</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>red</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
IRB name	interface <b>irb</b>
IRB units and addresses	logical unit 100: 192.0.2.1/25 logical unit 200: 192.0.2.129/25

For the sake of consistency and to avoid confusion, [Table 9 on page 42](#) shows IRB logical unit numbers that match the IDs of the corresponding VLANs. However, you do not have to assign logical unit numbers that match the VLAN IDs—you can use any values for the units. To bind the logical units of the IRB to the appropriate VLANs, you use the **l3-interface** statement.

Because IRBs operate at Layer 3, you can use Layer 3 services such as firewall filters or CoS rewriting with them.

[Table 10 on page 42](#) shows the number of IRBs/RVIs that each QFX platform supports.

Table 10: Number of Supported IRBs/RVIs by Platform

Platform	Number of Supported IRBs/RVIs
QFX5100	4000
EX4600	4000
OCX1100	4000
QFX3500	1200
QFX3000-G	1024
QFX3000-M	1024

**Related Documentation**

- [Example: Configuring Routing Between VLANs on One Switch on page 45](#)

## Configuring VLANs

Switches use VLANs to make logical groupings of network nodes with their own broadcast domains. You can use VLANs to limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.



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



**NOTE:** Two logical interfaces that are configured on the same physical interface cannot be mapped to the same VLAN.

For each endpoint on the VLAN, configure the following VLAN parameters on the corresponding interface:

1. Specify the description of the VLAN:

```
[edit interfaces interface-name unit 0]
user@switch# set description vlan-description
```

2. Specify the unique name of the VLAN:



**NOTE:** Switches that run Junos OS with the ELS configuration style do not support a default VLAN. Therefore, on such switches, you must explicitly configure at least one VLAN, even if your network is simple and you want only one broadcast domain to exist.

```
[edit interfaces interface-name unit 0]
user@switch# set family ethernet-switching vlan members vlan-name
```

3. Create the subnet for the VLAN:

```
[edit interfaces]
user@switch# set vlan unit 0 family inet address ip-address
```

4. Configure the VLAN tag ID or VLAN ID list for the VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id-number
or
```

```
[edit vlans]
user@switch# set vlan-name vlan-id-list [vlan-ids | vlan-id--vlan-id-]
```

### Related Documentation

- [Example: Setting Up Basic Bridging and a VLAN on the QFX Series](#)
- [Configuring IRB Interfaces on page 44](#)
- [Creating a Series of Tagged VLANs](#)
- [Understanding Bridging and VLANs](#)

## Configuring IRB Interfaces

Integrated routing and bridging (IRB) interfaces enable a switch to recognize which packets are being sent to local addresses so that they are bridged whenever possible and are routed only when needed. Whenever packets can be switched instead of routed, several layers of processing are eliminated. Switching also reduces the number of address look-ups.



**NOTE:** In versions of Junos OS that do not support Enhanced Layer 2 Software (ELS), this type of interface is called a routed VLAN interface (RVI).

To configure the routed VLAN interface:

1. Create the VLAN by assigning it a name and a VLAN ID:

```
[edit]
user@switch# set vlans support vlan-id 111
```

2. Assign an interface to the VLAN by specifying the logical interface (with the **unit** statement) and specifying the VLAN name as the member:

```
[edit]
user@switch# set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members
support
```

3. Create the subnet for the VLAN's broadcast domain:

```
[edit]
user@switch# set interfaces irb unit 111 family inet address 111.111.111.1/24
```

4. Bind a Layer 3 interface with the VLAN:

```
[edit]
user@switch# set vlans support l3-interface irb.111
```



**NOTE:** If you are using a version of Junos OS that does not support ELS, you create a Layer 3 virtual interface named **vlan**



**NOTE:** Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

You can display the configuration settings:

```
user@switch> show interfaces irb terse
Interface          Admin Link Proto  Local          Remote
vlan               up    up
irb.111            up    up  inet   111.111.111.1/24

user@switch> show vlans
Name      Tag      Interfaces
default                                     None
```

```

employee-vlan 20
                ge-1/0/0.0, ge-1/0/1.0, ge-1/0/2.0
marketing      40
                ge-1/0/10.0, ge-1/0/20.0, ge-1/0/30.0
support        111
                ge-0/0/18.0
mgmt
                bme0.32769, bme0.32771*

user@switch> show ethernet-switching table
Ethernet-switching table: 1 entries, 0 learned
  VLAN      MAC address      Type      Age Interfaces
  support    00:19:e2:50:95:a0 Static    - Router

```

**Related Documentation**

- [Understanding Integrated Routing and Bridging on page 41](#)

## Example: Configuring Routing Between VLANs on One Switch

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs). For example, you might want to create a VLAN that includes the employees in a department and the resources that they use often, such as printers, servers, and so on.

Of course, you also want to allow these employees to communicate with people and resources in other VLANs. To forward packets between VLANs you normally you need a router that connects the VLANs. However, you can accomplish this on a Juniper Networks switch without using a router by configuring an integrated routing and bridging (IRB) interface (also known as a routed VLAN interface—or RVI—in versions of Junos OS that do not support Enhanced Layer 2 Software). Using this approach reduces complexity and avoids the costs associated with purchasing, installing, managing, powering, and cooling another device.

- [Requirements on page 45](#)
- [Overview and Topology on page 45](#)
- [Configure Layer 2 switching for two VLANs on page 46](#)
- [Verification on page 49](#)

### Requirements


This example uses the following hardware and software components:

- One switch
- Junos OS Release 11.1 or later

### Overview and Topology

This example uses an IRB to route traffic between two VLANs on the same switch. The topology is shown in [Figure 1 on page 46](#).

Figure 1: IRB with One Switch



This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch and configuring an IRB to enable routing between the VLANs. One VLAN, called **blue**, is for the sales and marketing group, and a second, called **red**, is for the customer support team. The sales and support groups each have their own file servers and wireless access points. Each VLAN must have a unique name, tag (VLAN ID), and distinct IP subnet. [Table 11 on page 46](#) lists the components of the sample topology.

Table 11: Components of the Multiple VLAN Topology

Property	Settings
VLAN names and tag IDs	<b>blue</b> , ID 100 <b>red</b> , ID 200
Subnets associated with VLANs	<b>blue</b> : 192.0.2.0/25 (addresses 192.0.2.1 through 192.0.2.126) <b>red</b> : 192.0.2.128/25 (addresses 192.0.2.129 through 192.0.2.254)
Interfaces in VLAN <b>blue</b>	Sales server port: <b>xe-0/0/4</b> Sales wireless access points: <b>xe-0/0/6</b>
Interfaces in VLAN <b>red</b>	Support server port: <b>xe-0/0/0</b> Support wireless access points: <b>xe-0/0/2</b>
IRB name	interface <b>irb</b>
IRB units and addresses	logical unit 100: 192.0.2.1/25 logical unit 200: 192.0.2.129/25

This configuration example creates two IP subnets, one for the blue VLAN and the second for the red VLAN. The switch bridges traffic within the VLANs. For traffic passing between two VLANs, the switch routes the traffic using an IRB on which you have configured addresses in each IP subnet.

To keep the example simple, the configuration steps show only a few interfaces and VLANs. Use the same configuration procedure to add more interfaces and VLANs. By default, all interfaces are in access mode, so you do not have to configure the port mode.

### Configure Layer 2 switching for two VLANs

#### CLI Quick Configuration

To quickly configure Layer 2 switching for the two VLANs (**blue** and **red**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:



**NOTE:** The following example uses a version of Junos OS that supports Enhanced Layer 2 Software (ELS). When you use ELS, you create a Layer 3 virtual interface named **irb**. If you are using a version of Junos OS that does not support ELS, you create a Layer 3 virtual interface named **vlan**.

```
[edit]
set interfaces xe-0/0/4 unit 0 description "Sales server port"
set interfaces xe-0/0/4 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/6 unit 0 description "Sales wireless access point port"
set interfaces xe-0/0/6 unit 0 family ethernet-switching vlan members blue
set interfaces xe-0/0/0 unit 0 description "Support servers"
set interfaces xe-0/0/0 unit 0 family ethernet-switching vlan members red
set interfaces xe-0/0/2 unit 0 description "Support wireless access point port"
set interfaces xe-0/0/2 unit 0 family ethernet-switching vlan members red
set interfaces irb unit 100 family inet address 192.0.2.1/25
set interfaces irb unit 200 family inet address 192.0.2.129/25
set vlans blue l3-interface irb.100
set vlans blue vlan-id 100
set vlans red vlan-id 200
set vlans red l3-interface irb.200
```

#### Step-by-Step Procedure

To configure the switch interfaces and the VLANs to which they belong:

1. Configure the interface for the sales server in the blue VLAN:

```
[edit interfaces xe-0/0/4 unit 0]
user@switch# set description "Sales server port"
user@switch# set family ethernet-switching vlan members blue
```

2. Configure the interface for the wireless access point in the blue VLAN:

```
[edit interfaces xe-0/0/6 unit 0]
user@switch# set description "Sales wireless access point port"
user@switch# set family ethernet-switching vlan members blue
```

3. Configure the interface for the support server in the red VLAN:

```
[edit interfaces xe-0/0/0 unit 0]
user@switch# set description "Support server port"
user@switch# set family ethernet-switching vlan members red
```

4. Configure the interface for the wireless access point in the red VLAN:

```
[edit interfaces xe-0/0/2 unit 0]
user@switch# set description "Support wireless access point port"
user@switch# set family ethernet-switching vlan members red
```

#### Step-by-Step Procedure

Now create the VLANs and the IRB. The IRB will have logical units in the broadcast domains of both VLANs.

1. Create the red and blue VLANs by configuring the VLAN IDs for them:

```
[edit vlans]
user@switch# set blue vlan-id 100
user@switch# set red vlan-id 200
```

2. Create the interface named **irb** with a logical unit in the sales broadcast domain (blue VLAN):

```
[edit interfaces]
user@switch# set irb unit 100 family inet address 192.0.2.1/25
```

The unit number is arbitrary and does not have to match the VLAN tag ID. However, configuring the unit number to match the VLAN ID can help avoid confusion.

3. Add a logical unit in the support broadcast domain (red VLAN) to the **irb** interface:

```
[edit interfaces]
user@switch# set irb unit 200 family inet address 192.0.2.129/25
```

4. Complete the IRB configuration by binding the red and blue VLANs (Layer 2) with the appropriate logical units of the **irb** interface (Layer 3):

```
[edit vlans]
user@switch# set blue l3-interface irb.100
user@switch# set red l3-interface irb.200
```

Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  xe-0/0/4 {
    unit 0 {
      description "Sales server port";
      family ethernet-switching {
        vlan members blue;
      }
    }
  }
  xe-0/0/6 {
    unit 0 {
      description "Sales wireless access point port";
      family ethernet-switching {
        vlan members blue;
      }
    }
  }
  xe-0/0/0 {
    unit 0 {
      description "Support server port";
      family ethernet-switching {
        vlan members red;
      }
    }
  }
  xe-0/0/2 {
    unit 0 {
      description "Support wireless access point port";
      family ethernet-switching {
        vlan members red;
      }
    }
  }
  irb {
    unit 100 {
      family inet address 192.0.2.1/25;
    }
    unit 200 {
      family inet address 192.0.2.129/25;
    }
  }
}
```



```

    }
  }
}
vllans {
  blue {
    vlan-id 100;
    interface xe-0/0/4.0;
    interface xe-0/0/6.0;
    l3-interface irb 100;
  }
  red {
    vlan-id 200;
    interface xe-0/0/0.0;
    interface xe-0/0/2.0;
    l3-interface irb 200;
  }
}
}

```



**TIP:** To quickly configure the blue and red VLAN interfaces, issue the **load merge terminal** command, copy the hierarchy, and paste it into the switch terminal window.

## Verification

To verify that the **blue** and **red** VLANs have been created and are operating properly, perform these tasks:

- [Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces on page 49](#)
- [Verifying That Traffic Can Be Routed Between the Two VLANs on page 50](#)

### Verifying That the VLANs Have Been Created and Associated with the Correct Interfaces

**Purpose** Verify that the VLANs **blue** and **red** have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

**Action** List all VLANs configured on the switch:

```

user@switch> show vlans
Name      Tag      Interfaces
default   100      xe-0/0/0.0, xe-0/0/2.0, xe-0/0/4.0, xe-0/0/6.0,
blue      100      xe-0/0/4.0, xe-0/0/6,
red       200      xe-0/0/0.0, xe-0/0/2.0, *
mgmt      me0.0*

```

**Meaning** The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **blue** and **red** VLANs have been created. The **blue** VLAN has a tag ID of 100 and is associated with interfaces

**xe-0/0/4.0** and **xe-0/0/6.0**. VLAN **red** has a tag ID of 200 and is associated with interfaces **xe-0/0/0.0** and **xe-0/0/2.0**.

### Verifying That Traffic Can Be Routed Between the Two VLANs

**Purpose** Verify routing between the two VLANs.

**Action** Verify that the IRB logical units are up:

```
user@switch> show interfaces terse
irb.100                up    up    inet    192.0.2.1/25
irb.200                up    up    inet    192.0.2.129/25
```



**NOTE:** At least one port (access or trunk) with an appropriate VLAN assigned to it must be up for the irb interface to be up.

Verify that switch has created routes that use the IRB logical units:

```
user@switch> show route
192.0.2.0/25           *[Direct/0] 1d 03:26:45
                      > via irb.100
192.0.2.1/32           *[Local/0] 1d 03:26:45
                      Local via irb.100
192.0.2.128/25         *[Direct/0] 1d 03:26:45
                      > via irb.200
192.0.2.129/32         *[Local/0] 1d 03:26:45
                      Local via irb.200
```

List the Layer 3 routes in the switch's Address Resolution Protocol (ARP) table:

```
user@switch> show arp
MAC Address           Address           Name              Flags
00:00:0c:06:2c:0d     192.0.2.7        irb.100           None
00:13:e2:50:62:e0     192.0.2.132      irb.200           None
```

**Meaning** The output of the **show interfaces** and **show route** commands show that the Layer 3 IRB logical units are working and the switch has used them to create direct routes that it will use to forward traffic between the VLAN subnets. The **show arp** command displays the mappings between the IP addresses and MAC addresses for devices on both **irb.100** (associated with VLAN **blue**) and **irb.200** (associated with VLAN **red**). These two devices can communicate.

**Related Documentation**

- [Understanding Integrated Routing and Bridging on page 41](#)
- [irb \(Interfaces\) on page 108](#)
- [I3-interface on page 111](#)

## PART 5

# Configuration Statements and Operational Commands

- [Configuration Statements for Interfaces on page 53](#)
- [Configuration Statements for GRE on page 95](#)
- [Configuration Statements for LAGs and LACP on page 97](#)
- [Configuration Statements for IRB Interfaces on page 107](#)
- [Interfaces Operational Commands on page 113](#)
- [LAGs and LACP Operational Commands on page 217](#)



## CHAPTER 5

# Configuration Statements for Interfaces

- [address on page 54](#)
- [alarm \(chassis\) on page 57](#)
- [configured-flow-control on page 58](#)
- [description \(Interfaces\) on page 59](#)
- [ether-options on page 60](#)
- [eui-64 on page 61](#)
- [family on page 62](#)
- [flow-control on page 66](#)
- [fpc on page 67](#)
- [hold-time \(Physical Interface\) on page 68](#)
- [inet \(interfaces\) on page 69](#)
- [inet6 \(interfaces\) on page 70](#)
- [interface-range on page 71](#)
- [interfaces on page 73](#)
- [link-mode on page 80](#)
- [link-speed on page 81](#)
- [loopback \(Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet\) on page 82](#)
- [member on page 82](#)
- [member-range on page 83](#)
- [mtu on page 84](#)
- [pic on page 86](#)
- [rx-buffers on page 87](#)
- [source on page 88](#)
- [targeted-broadcast on page 89](#)
- [traceoptions \(Individual Interfaces\) on page 90](#)
- [tx-buffers on page 91](#)
- [unit on page 93](#)

## address

```

Syntax  address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        master-only;
        multipoint-destination address dlcid dlcid-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
                 length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        (vrrp-group | vrrp-inet6-group) group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-type authentication;
            authentication-key key;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority-number number;
            track {
                priority-cost seconds;
                priority-hold-time interface-name {
                    interface priority;
                    bandwidth-threshold bits-per-second {
                        priority;
                    }
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}

```

**Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family *family*],  
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*  
 family *family*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
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 the interface address.

**Options** *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.



**NOTE:** If you configure the same address on multiple interfaces in the same routing instance, Junos OS uses only the first configuration, the remaining address configurations are ignored and can leave interfaces without an address. Interfaces that do not have an assigned address cannot be used as a donor interface for an unnumbered Ethernet interface.

For example, in the following configuration the address configuration of interface xe-0/0/1.0 is ignored:

```
interfaces {
  xe-0/0/0 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.1.1/24;
      }
    }
  }
}
```

For more information on configuring the same address on multiple interfaces, see [“Configuring the Interface Address” on page 8](#).

The remaining statements are explained separately.



**NOTE:** The `edit logical-systems` hierarchy is not available on QFabric systems.

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


- Related Documentation**
- *Configuring the Protocol Family*
  - *Junos OS Administration Library for Routing Devices*
  - *family*
  - *negotiate-address*
  - *unnumbered-address (Ethernet)*



## alarm (chassis)

<b>Syntax</b>	<pre>alarm {     interface-type {         alarm-name (ignore   red   yellow);     } }</pre>
<b>Hierarchy Level</b>	<p>[edit chassis],          [edit chassis interconnect-device <i>name</i>],          [edit chassis node-group <i>name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.          Statement introduced in Junos OS Release 11.1 for the QFX Series.          Statement introduced in Junos OS Release 12.2 for the ACX Series.          Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure the chassis alarms and whether they trigger a red or yellow alarm, or whether they are ignored. Red alarm conditions light the <b>RED ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the contact on the craft interface or LCD screen. Yellow alarm conditions light the <b>YELLOW ALARM</b> LED on either the router's craft interface or the switch's LCD screen and trigger an audible alarm if one is connected to the craft interface or LCD screen.</p> <p>To configure more than one alarm, include multiple <i>alarm-name</i> lines.</p>
<b>Options</b>	<p><i>alarm-name</i>—Alarm condition. For a list of conditions, see <i>System-Wide Alarms and Alarms for Each Interface Type</i>.</p> <p><i>ignore</i>—The specified alarm condition does not set off any alarm.</p> <p><i>interface-type</i>—Type of interface on which you are configuring the alarm: <b>atm</b>, <b>ethernet</b>, <b>sonet</b>, or <b>t3</b>.</p> <p><b>red</b>—The specified alarm condition sets off a red alarm.</p> <p><b>yellow</b>—The specified alarm condition sets off a yellow alarm.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.          interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Understanding Alarms</i></li> <li>• <i>Chassis Conditions That Trigger Alarms</i></li> <li>• <i>Chassis Alarm Messages on a QFX3500 Device</i></li> <li>• <i>Interface Alarm Messages</i></li> </ul>

## configured-flow-control

<b>Syntax</b>	<pre>configured-flow-control {     rx-buffers (on   off);     tx-buffers (on   off); }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the <b>rx-buffers</b> and the <b>tx-buffers</b> values when you configure asymmetric flow control.</p> <p>Use the <b>flow-control</b> and <b>no-flow-control</b> statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>
<b>Default</b>	Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">congestion-notification-profile</a></li> <li><a href="#">flow-control on page 66</a></li> <li><i>Configuring CoS Asymmetric Ethernet PAUSE Flow Control</i></li> <li><i>Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control</i></li> <li><i>Understanding CoS Flow Control (Ethernet PAUSE and PFC)</i></li> </ul>

## description (Interfaces)

<b>Syntax</b>	<code>description text;</code>
<b>Hierarchy Level</b>	<code>[edit <a href="#">interfaces</a> interface-name],</code> <code>[edit <a href="#">interfaces</a> interface-name unit logical-unit-number],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]</code>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the <b>show interfaces</b> commands, and is also exposed in the <b>ifAlias</b> Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p>
<b>Options</b>	<b>text</b> —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Interface Description</i></li> <li>• <i>Adding a Logical Unit Description to the Configuration</i></li> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li> <li>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i></li> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> <li>• <i>Using DHCP Relay Agent Option 82 Information</i></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> <li>• <i>Example: Connecting Access Switches to a Distribution Switch</i></li> </ul>

## ether-options

**Syntax** The **auto-negotiation** and **speed** statements are not supported on the OCX Series.

```
ether-options {
  802.3ad aex {
    lacp {
      force-up;
      (primary | backup);
    }
  }
  (auto-negotiation | no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  ethernet-switch-profile
    storm-control storm-control-profile;
}
(flow-control | no-flow-control);
link-mode mode;
(loopback | no-loopback);
speed (auto-negotiation | no-auto-negotiation);
}
```

**Hierarchy Level** [edit [interfaces](#) *interface-name*]

**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 **ether-options** properties for a Gigabit Ethernet or 10-Gigabit Ethernet interface.



**NOTE:** The **auto-negotiation** and **speed** statements are not supported on the OCX Series.

The statements are explained separately.

**Default** Enabled.

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

**Related Documentation**

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

## eui-64

---

<b>Syntax</b>	eui-64;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Interface Address on page 8</a></li></ul>

## family

**Syntax** The **ethernet-switching** statement and all of its substatements are not supported on OCX Series switches.

```
family {
  ethernet-switching {
    filter {
      group filter-group-number;
      input filter-name;
      input-list [ filter-names ];
      output filter-name;
      output-list [ filter-names ];
    }
    interface-mode (access | trunk);
    recovery-timeout seconds;
    storm-control profile-name;
    vlan {
      members (vlan-name [ -vlan-names ] | all);
    }
  }
}
inet {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
}
address ipv4-address {
  arp ip-address (mac | multicast-mac) mac-address <publish>;
  broadcast address;
  preferred;
  primary;
  vrrp-group group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    advertisements-threshold number;
    authentication-key key;
    authentication-type authentication;
    fast-interval milliseconds;
    (preempt | no-preempt) {
      hold-time seconds;
    }
  }
  priority number;
  track {
    interface interface-name {
      priority-cost number;
    }
  }
  priority-hold-time seconds;
  route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-address [addresses];
vrrp-inherit-from {
```

```

        active-group group-number;
        active-interface interface-name;
    }
}
filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
}

```

```

        vrrp-inherit-from {
            active-group group-name;
            active-interface interface-name;
        }
    }
    (dad-disable | no-dad-disable);
    filter {
        group filter-group-number;
        input filter-name;
        input-list [ filter-names ];
        output filter-name;
        output-list [ filter-names ];
    }
    mtu bytes;
    nd6-stale-time time;
    no-neighbor-learn;
    no-redirects;
    policer {
        input policer-name;
        output policer-name;
    }
    rpf-check {
        fail-filter filter-name;
        mode {
            loose;
        }
    }
    }
    mpls {
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        mtu bytes;
    }
}

```

<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ], [edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure protocol family information for the logical interface on the QFX Series and OCX Series product.



**Default**

**NOTE:** The **ethernet-switching** statement and all of its substatements are not supported on OCX Series switches.

Access interfaces on the QFX Series are set to **family ethernet-switching** by default. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting first.

You must configure a logical interface to be able to use the physical device.

**Options**

Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)



**NOTE:** Routed VLAN interfaces, also referred to as integrated routing and bridging (IRB) interfaces, are not supported on OCX Series switches.

- 10-Gigabit Ethernet (**xe**)

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

**Required Privilege Level**


interface—To view this statement in the configuration.

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

**Related Documentation**

- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13](#)
- [Configuring Link Aggregation on page 33](#)
- [Configuring IRB Interfaces on page 44](#)
- *Junos OS Network Interfaces Library for Routing Devices*

## flow-control


<b>Syntax</b>	(flow-control   no-flow-control);
<b>Hierarchy Level</b>	[edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit <b>interfaces</b> <i>interface-name</i> <b>ether-options</b> hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> <p>OCX Series switches do not support PFC.</p> </div> <ul style="list-style-type: none"> <li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li> <li>• <b>no-flow-control</b>—Disable flow control.</li> </ul>
<b>Default</b>	Flow control is disabled.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">configured-flow-control on page 58</a></li> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> </ul>

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\)](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

## fpc

<b>Syntax</b>	<pre>fpc slot {   auto-speed-detection disable;   pic <i>pic-number</i> {     <b>tunnel-port</b> <i>port-number</i> tunnel-services;     port <i>port-number</i> {       channel-speed (<i>speed</i> disable-auto-speed-detection) ;     }     port-range <i>port-range-low port-range-high</i> {       channel-speed (<i>speed</i> disable-auto-speed-detection);     }   } }</pre>
<b>Hierarchy Level</b>	[edit chassis]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.</p> <p>For generic routing encapsulation (GRE) tunneling, use the <b>tunnel-port</b> statement to specify the port that you want to convert to a GRE tunnel port.</p>
<b>Options</b>	<p><b>slot</b>—Number of the FPC slot. For QFX3500, QFX3600, and OCX Series devices, the slot number is always 0.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show chassis fpc</a></li> <li>• <a href="#">Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</a></li> </ul>

## hold-time (Physical Interface)

<b>Syntax</b>	<code>hold-time up <i>milliseconds</i> down <i>milliseconds</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> ], [edit <a href="#">interfaces</a> <i>interface-range</i> <i>interface-range-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 10.4R5 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the <b>hold-time</b> value to use to damp shorter interface transitions milliseconds. When an interface goes from up to down, it is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly, an interface is not advertised as being up until it has remained up for the hold-time period.
	<div>  <b>NOTE:</b> <ul style="list-style-type: none"> <li>We recommend that you configure the hold-time value after determining an appropriate value by performing repeated tests in the actual hardware environment. This is because the appropriate value for hold-time depends on the hardware (XFP, SFP, SR, ER, or LR) used in the networking environment.</li> <li>The hold-time option is not available for controller interfaces.</li> </ul> </div>
<b>Default</b>	Interface transitions are not damped.
<b>Options</b>	<p><b>down <i>milliseconds</i></b>—Hold time to use when an interface transitions from up to down. Junos OS advertises the transition within 100 milliseconds of the time value you specify.</p> <p><b>Range:</b> 0 through 4,294,967,295</p> <p><b>Default:</b> 0 (interface transitions are not damped)</p> <p><b>up <i>milliseconds</i></b>—Hold time to use when an interface transitions from down to up. Junos OS advertises the transition within 100 milliseconds of the time value you specify.</p> <p><b>Range:</b> 0 through 4,294,967,295</p> <p><b>Default:</b> 0 (interface transitions are not damped)</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><code>advertise-interval</code></li> <li><code>interfaces</code> (for EX Series switches)</li> </ul>

- *Physical Interface Damping Overview*
- *Damping Shorter Physical Interface Transitions*
- *Damping Longer Physical Interface Transitions*

## inet (interfaces)

<b>Syntax</b>	<pre>inet {   address <i>address</i> {     primary;     filter input <i>filter-name</i>;     filter output <i>filter-name</i>;     targeted-broadcast;   } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family], [edit <a href="#">interfaces</a> <a href="#">interface-range</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the primary IP address for the logical interface.
<b>Default</b>	You must configure a logical interface to be able to use the physical device.
<b>Options</b>	The remaining statements are explained separately.—
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> </ul>

## inet6 (interfaces)

---

<b>Syntax</b>	<pre>inet6 {     address address {         eui-64         preferred         primary;         filter input <i>filter-name</i>;         filter output <i>filter-name</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces interface-name unit logical-unit-number</a> family], [edit <a href="#">interfaces interface-range interface-name unit logical-unit-number</a> family]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the primary IP address for the logical interface.
<b>Default</b>	You must configure a logical interface to be able to use the physical device.
<b>Options</b>	The remaining statements are explained separately.—
<b>Required Privilege Level</b>	interface—To view this statement in the configuration.interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li></ul>

## interface-range

**Syntax** The `vlan-id` statement is not supported on OCX Series switches.

```
interface-range interface-range-name {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
    (auto-negotiation| no-auto-negotiation);
    (flow-control | no-flow-control);
    link-mode mode;
    speed (auto-negotiation | speed);
  }
  hold-time milliseconds down milliseconds;
  member interface-name;
  member-range starting-interface-name to ending-interface-name;
  mtu bytes;
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    (traps | no traps);
    vlan-id vlan-id-number;
  }
}
```

**Hierarchy Level** [edit [interfaces](#)]

**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**



**NOTE:** The `vlan-id` statement and Fibre Channel interfaces are not supported on OCX Series switches.

Group interfaces that share a common configuration profile.



**NOTE:** The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces.

**Options** *interface-range-name*—Name of the interface range.



**NOTE:** You can use regular expressions and wildcards to specify the interfaces in the member range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

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

**Related Documentation**

- [Understanding Interface Ranges on page 5](#)
- [Interfaces Overview](#)
- [Interfaces Overview on page 3](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)



## interfaces

**Syntax** The following statements and their associated substatements are not supported on OCX Series switches: **auto-negotiation**, **speed**, **ethernet-switching**, **fcoe-lag**, **fibre-channel**, **fibrechannel-options**, **mc-ae**, **vlan**, **vlan-id**, and **vlan-tagging**.

```

interfaces {
  aex {
    disable;
    aggregated-ether-options {
      configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
      }
      (fcoe-lag | no-fcoe-lag);
      (flow-control | no-flow-control);
      lacp mode {
        admin-key key;
        force-up;
        periodic interval;
        system-id mac-address;
      }
      link-speed speed;
      local-bias (edit interfaces ae);
      loopback;
      no-loopback;
      minimum-links number;
    }
    mc-ae {
      chassis-id chassis-id;
      mc-ae-id mc-ae-id;
      mode (active-active);
      status-control (active | standby);
    }
    description text;
    gratuitous-arp-reply | no-gratuitous-arp-reply)
    hold-time down milliseconds up milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    short-reach-mode (enable | disable);
    traceoptions;
    (traps | no traps);
    unit logical-unit-number {
      disable;
      description text;
      family {
        ethernet-switching {
          filter input filter-name;
          filter output filter-name;
          native-vlan-id vlan-id;
          port-mode mode;
          reflective-relay;
          vlan {
            members [ (all | names | vlan-ids) ];
          }
        }
      }
    }
  }
}

```

```

    }
  }
  inet {
    address address {
      primary;
    }
    filter input filter-name;
    filter output filter-name;
    primary;
    targeted-broadcast;
  }
  (traps | no traps);
  vlan-id vlan-id-number;
}
vlan-tagging;
}
interface-range interface-range-name {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
  }
  (auto-negotiation | no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  (flow-control | no-flow-control);
  link-mode mode;
  speed (auto-negotiation | speed);
}
hold-time milliseconds down milliseconds;
member interface-name;
member-range starting-interface-name to ending-interface-name;
mtu bytes;
unit logical-unit-number {
  disable;
  description text;
  family family-name {...}
  (traps | no traps);
  vlan-id vlan-id-number;
}
}
lo0 {
  disable;
  description text;
  hold-time milliseconds down milliseconds;
  traceoptions;
  (traps | no traps);
  unit logical-unit-number {
    disable;
    description text;
    family {

```

```

    inet {
        address address {
            primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
    }
    (traps | no traps);
}
}
mex {
    disable;
    description text;
    hold-time milliseconds down milliseconds;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    no-gratuitous-arp-request;
    traceoptions;
    traps;
    unit logical-unit-number {
        disable;
        description text;
        family {
            ethernet-switching {
                filter input filter-name;
                filter output filter-name;
                native-vlan-id vlan-id;
                port-mode mode;
                reflective-relay;
                vlan {
                    members [ (all | names | vlan-ids) ];
                }
            }
        }
        inet {
            address address {
                primary;
                filter input filter-name;
                filter output filter-name;
                primary;
                targeted-broadcast;
            }
        }
        traps;
        vlan-id vlan-id-number;
    }
}
vlan-tagging;
vlan {
    disable;
    description text;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no traps);
}

```

```

unit logical-unit-number {
  description text;
  disable;
  family {
    inet {
      address address {
        primary;
      }
      filter input filter-name;
      filter output filter-name;
      primary;
      targeted-broadcast;
    }
    (traps | no traps);
  }
}
fc-0/0/port {
  fibrechannel-options {
    bb-sc-n;
    (loopback | no-loopback);
    speed (auto-negotiation | 2g | 4g | 8g);
  }
  unit logical-unit-number {
    disable;
    description text;
    family {
      fibre-channel {
        port-mode np-port;
      }
    }
    (traps | no traps);
  }
}
ge-0/0/port {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
        primary;
      }
    }
  }
  (auto-negotiation | no-auto-negotiation);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  (flow-control | no-flow-control);
  link-mode mode;
  loopback;
  no-loopback;
  speed (auto-negotiation | speed);
}
gratuitous-arp-reply| no-gratuitous-arp-reply);
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;

```

```

traceoptions;
(traps | no traps);
unit logical-unit-number {
  description text;
  disable;
  family {
    ethernet-switching {
      filter input filter-name;
      filter output filter-name;
      native-vlan-id vlan-id;
      port-mode mode;
      reflective-relay;
      vlan {
        members [ (all | names | vlan-ids) ];
      }
    }
  }
  inet {
    address address {
      primary;
    }
    filter input filter-name;
    filter output filter-name;
    primary;
    targeted-broadcast;
  }
  (traps | no traps);
  vlan-id vlan-id-number;
}
vlan-tagging;
}
vrrp-group group-id {
  (accept-data | no-accept-data);
  advertise-interval seconds;
  authentication-key key;
  authentication-type authentication;
  fast-interval milliseconds;
  (preempt | no-preempt) {
    hold-time seconds;
  }
  priority number;
  track {
    interface interface-name {
      bandwidth-threshold bits-per-second priority-cost priority;
      priority-cost priority;
    }
    priority-hold-time seconds;
    route prefix/prefix-length routing-instance instance-name priority-cost priority;
  }
}
virtual-address [ addresses ];
}
xe-0/0/port {
  disable;
  description text;
  ether-options {
    802.3ad aex {

```

```

    lacp {
        force-up;
        (primary | backup);
    }
}
configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
}
(flow-control | no-flow-control);
loopback;
no-loopback;
}
(gratuitous-arp-reply | no-gratuitous-arp-reply)
hold-time milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
traceoptions;
(traps | no traps);
unit logical-unit-number {
    disable;
    description text;
    family {
        ethernet-switching {
            filter input filter-name;
            filter output filter-name;
            native-vlan-id vlan-id;
            port-mode mode;
            reflective-relay;
            vlan {
                members [ (all | names | vlan-ids) ];
            }
        }
        fibre-channel {
            port-mode (f-port | np-port);
        }
    }
    inet {
        address address {
            primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
    }
    (traps | no traps);
    vlan-id vlan-id-number;
}
vlan-tagging;
}
}

```

Hierarchy Level [edit]

<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Configure the interfaces on the QFX Series and OCX Series.</p> <p>The following statements and their associated substatements are not supported on OCX Series switches: <b>auto-negotiation</b>, <b>ethernet-switching</b>, <b>fcoe-lag</b>, <b>fibre-channel</b>, <b>fibrechannel-options</b>, <b>mc-ae</b>, <b>speed</b>, <b>vlan</b>, <b>vlan-id</b>, and <b>vlan-tagging</b></p> <p>Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.</p>
<b>Options</b>	<p><b>aex</b>—Configure an aggregated Ethernet interface.</p> <p><b>xe-0/0/</b><i>port</i><b>/</b>—Configure a 10-Gigabit Ethernet interface.</p> <p><b>ge-0/0/</b><i>port</i><b>/</b>—Configure a Gigabit Ethernet interface.</p> <p><b>fc-0/0/</b><i>port</i><b>/</b>—Configure a Fibre Channel interface.</p> <p><b>meX</b>/—Configure a management interface.</p> <p><b>mc-ae</b>—Configure a multichassis aggregated Ethernet (MC-AE) interface.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Interfaces Overview</i></li> <li>• <a href="#">Understanding Interface Ranges on page 5</a></li> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li> <li>• <a href="#">Configuring Link Aggregation on page 33</a></li> <li>• <i>Configuring a Layer 3 Logical Interface</i></li> </ul>


## link-mode

---

<b>Syntax</b>	<code>link-mode mode;</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces interface-name ether-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Set the device's link-connection characteristic.
<b>Default</b>	The <b>full-duplex</b> mode is enabled.
<b>Options</b>	<p><b>mode</b> —Link characteristic:</p> <ul style="list-style-type: none"><li>• <b>full-duplex</b>—Connection is full duplex.</li><li>• <b>half-duplex</b>—Connection is half duplex.</li><li>• <b>automatic</b>—Link mode is negotiated.</li></ul> <p>If <b>no-auto-negotiation</b> is specified in the <b>ether-options</b> option, you can select only <b>full-duplex</b> or <b>half-duplex</b>. If <b>auto-negotiation</b> is specified in the <b>ether-options</b> option, you can select any mode.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li><li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li></ul>



## link-speed

<b>Syntax</b>	link-speed <i>speed</i> ;
<b>Hierarchy Level</b>	[edit interfaces aex <a href="#">aggregated-ether-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	For aggregated Ethernet interfaces only, set the required link speed.
<b>Options</b>	<p><b>speed</b>—For aggregated Ethernet links, you can specify the speed in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p>On QFX5100 and EX4600 standalone switches and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.</p> <p>Aggregated Ethernet links on the QFX Series can have one of the following speed values:</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> <b>NOTE:</b> OCX Series switches only support 10g and 40g interfaces. Mixed rate aggregated Ethernet interfaces are not support on the OCX Series.</p> </div> <ul style="list-style-type: none"> <li>• <b>100g</b>—Links are 100 Gbps.</li> <li>• <b>100m</b>—Links are 100 Mbps.</li> <li>• <b>10g</b>—Links are 10 Gbps.</li> <li>• <b>1g</b>—Links are 1 Gbps.</li> <li>• <b>40g</b>—Links are 40 Gbps.</li> <li>• <b>50g</b>—Links are 50 Gbps.</li> <li>• <b>80g</b>—Links are 80 Gbps.</li> <li>• <b>8g</b>—Links are 8 Gbps.</li> <li>• <b>0c192</b>—Links are OC-192.</li> <li>• <b>mixed</b>—Links are 10 Gbps and 40Gbps.</li> </ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 33</a></li> </ul>

## loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet)

---

<b>Syntax</b>	(loopback   no-loopback);
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	For aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Ethernet Loopback Capability on page 14</a></li></ul>

## member

---



<b>Syntax</b>	member <i>interface-name</i> ;
<b>Hierarchy Level</b>	[edit <a href="#">interfaces interface-range</a> <i>interface-range-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the name of the member interface belonging to an interface range on the QFX Series switch.
<b>Options</b>	<i>interface-name</i> —Name of the interface.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces</a></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li><li>• <a href="#">Interfaces Overview</a></li><li>• <a href="#">Interfaces Overview on page 3</a></li><li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li></ul>

## member-range

---

<b>Syntax</b>	<code>member-range <i>starting-interface-name ending-interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-range interface-range-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.
<b>Options</b>	<i>starting interface-name ending interface-name</i> —Name of the first member and the name of the last member in the interface sequence.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Understanding Interface Ranges on page 5</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces</a></li> <li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li> <li>• <a href="#">Interfaces Overview</a></li> <li>• <a href="#">Interfaces Overview on page 3</a></li> <li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li> </ul>

## mtu

<b>Syntax</b>	<code>mtu bytes;</code>
<b>Hierarchy Level</b>	[edit <code>interfaces interface-name</code> ], [edit <code>interfaces interface-range interface-name</code> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On QFX3500, QFX3600, QFX5100, and OCX Series switches, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight.</p> <p>Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:</p> <ul style="list-style-type: none"> <li>• <b>For LAG interfaces</b>—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (<code>aex</code>) automatically configures the jumbo MTU size on the member links.</li> <li>• <b>For RVIs</b>—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named <code>vlan</code>. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the <code>vlan</code> interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named <code>vlan</code> (the RVI). On a QFX5100 switch jumbo frames on the RVI are configured on the basis of the interface MTU.</li> </ul> <div style="margin-top: 20px;"> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE:</b> RVIs are not supported on OCX Series switches.</p> </div> </div> <div style="margin-top: 20px;"> <div style="display: flex; align-items: center;">  <div> <p><b>CAUTION:</b> Setting or deleting the jumbo MTU size on the RVI (the <code>vlan</code> interface) while the switch is transmitting packets might result in dropped packets.</p> </div> </div> </div> </div>
<b>Options</b>	<p><code>bytes</code> —MTU size.</p> <p><b>Range:</b> 64 through 9216 bytes</p>

**Default:** 1514 bytes

<b>Required Privilege</b>	interface—To view this statement in the configuration.
<b>Level</b>	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## pic

**Syntax** `pic pic-number {  
     tunnel-port port-number tunnel-services;  
     port port-number {  
         channel-speed (speed|disable-auto-speed-detection) ;  
     }  
     port-range port-range-low port-range-high {  
         channel-speed (speed|disable-auto-speed-detection) ;  
     }  
 }`

**Hierarchy Level** [edit chassis fpc *slot*]

**Release Information** Option **channel-speed** introduced in Junos OS Release 13.2 for the QFX Series.



**NOTE:** This statement is not supported on the OCX Series.

**Description** (QFX3500, QFX3600, and QFX5100 standalone switches running Enhanced Layer 2 Software only)—Configure a specific port or a range of ports to operate as 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.

**Options** **pic *pic-number***—(QFX3500 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **1** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.  
 (QFX3600 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.

**port *physical-port-number***—Port number on which you want to configure the port type.

**port-range-low**—Lowest-numbered port in the range of ports.

**port-range-high**—Highest-numbered port in the range of ports.


**channel-speed (*speed* |disable-auto-speed-detection)** —Configure *10g* for 10-Gigabit Ethernet type ports, and configure *disable-auto-speed-detection* to disable auto-channelization.

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

**Related Documentation**

- *Channelizing Interfaces*

## rx-buffers

<b>Syntax</b>	rx-buffers (on   off);
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>rx-buffers</b> statement with the <b>tx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p>
	<div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>
<b>Default</b>	Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.
<b>Options</b>	<b>on   off</b> —Enable or disable an interface to generate and send Ethernet PAUSE messages.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 66</a></li> <li>• <a href="#">tx-buffers on page 91</a></li> </ul>

- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*

---

## source

<b>Syntax</b>	<code>source source-address;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the source address of the tunnel.
<b>Default</b>	If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.
<b>Options</b>	<b><i>source-address</i></b> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li></ul>




## targeted-broadcast

---


<b>Syntax</b>	targeted-broadcast;
<b>Hierarchy Level</b>	[edit <b>interfaces</b> <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> family inet], [edit <b>interfaces</b> <i>interface-range</i> <i>interface-range-name</i> <b>unit</b> <i>logical-unit-number</i> family inet]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify whether the IP packets destined for a Layer 3 broadcast need to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.
<b>Default</b>	When this statement is not included, broadcast packets are sent to the Routing Engine only.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring IP Directed Broadcast on an EX Series Switch</i></li> <li>• <i>Configuring IP Directed Broadcast (CLI Procedure)</i></li> <li>• <i>Understanding IP Directed Broadcast for EX Series Switches</i></li> </ul>

## traceoptions (Individual Interfaces)

---

<b>Syntax</b>	<pre>traceoptions {     flag <i>flag</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p>The <b>traceoptions</b> statement for interfaces does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system <b>syslog</b> file in the directory <b>/var/log</b>.</p> <div><b>NOTE:</b> The traceoptions statement is not supported on the QFX3000 QFabric system.</div>
<b>Default</b>	If you do not include this statement, no interface-specific tracing operations are performed.
<b>Options</b>	<p><b>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements. The following are the interface-specific tracing options.</p> <ul style="list-style-type: none"><li>• <b>all</b>—All interface tracing operations</li><li>• <b>event</b>—Interface events</li><li>• <b>ipc</b>—Interface interprocess communication (IPC) messages</li><li>• <b>media</b>—Interface media changes</li><li>• <b>q921</b>—ISDN Q.921 frames</li><li>• <b>q931</b>—ISDN Q.931 frames</li></ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Tracing Operations of an Individual Router or Switch Interface</i></li></ul>

## tx-buffers

<b>Syntax</b>	tx-buffers (on   off);
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> <a href="#">configured-flow-control</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for the QFX Series.
<b>Description</b>	<p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>tx-buffers</b> statement with the <b>rx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p>
	<div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>
<b>Default</b>	Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.
<b>Options</b>	<b>on   off</b> —Enable or disable an interface to respond to an Ethernet PAUSE message.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 66</a></li> <li>• <a href="#">rx-buffers on page 87</a></li> </ul>

- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*

## unit

**Syntax** The **ethernet-switching** and **fibre-channel** statements and all of their substatements are not supported on OCX Series switches.

```
unit logical-unit-number {
  family {
    ethernet-switching {
      filter input filter-name;
      filter output filter-name;
      native-vlan-id vlan-id;
      port-mode mode;
      vlan {
        members [ (all | names | vlan-ids) ];
      }
    }
    fibre-channel {
      port-mode (f-port | np-port);
    }
    inet {
      address address {
        primary;
      }
      filter input filter-name;
      filter output filter-name;
      primary;
      targeted-broadcast;
    }
  }
}
```

**Hierarchy Level** [edit **interfaces** *interface-name*],  
[edit **interfaces** **interface-range** *interface-range-name*]

**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**



**NOTE:** The **ethernet-switching** and **fibre-channel** statements and all of their substatements are not supported on OCX Series switches.

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

**Default** You must configure a logical interface to be able to use the physical device.

**Options** *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 16,384

The remaining statements are explained separately.

<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li><li>• <a href="#">Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 13</a></li><li>• <a href="#">Configuring Link Aggregation on page 33</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## CHAPTER 6

# Configuration Statements for GRE

- [destination \(Tunnels\) on page 95](#)
- [tunnel on page 96](#)
- [tunnel-port on page 96](#)

## destination (Tunnels)

---

<b>Syntax</b>	<code>destination address;</code>
<b>Hierarchy Level</b>	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address</code> <code>    <i>interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>    family inet address <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>    family inet unnumbered-address <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code>    tunnel]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.
<b>Options</b>	<b><i>address</i></b> —Address of the remote side of the connection.
<b>Required Privilege Level</b>	<b>interface</b> —To view this statement in the configuration. <b>interface-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Interface Address on page 8</a></li><li>• <a href="#">point-to-point</a></li></ul>

## tunnel

---

<b>Syntax</b>	<pre>tunnel {     destination destination-address;     source source-address;     ttl number; (not supported on QFX and OCX Series switches) }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	<p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li></ul>

## tunnel-port

---

<b>Syntax</b>	<pre>tunnel-port <i>port-number</i> tunnel-services;</pre>
<b>Hierarchy Level</b>	[edit chassis fpc <i>slot</i> pic <i>pic-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the port number for generic routing encapsulation (GRE) tunneling.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i></li></ul>



## CHAPTER 7

# Configuration Statements for LAGs and LACP

- [aggregated-devices on page 98](#)
- [aggregated-ether-options on page 99](#)
- [chassis on page 101](#)
- [802.3ad on page 102](#)
- [device-count on page 103](#)
- [ethernet on page 103](#)
- [force-up on page 104](#)
- [lacp \(802.3ad\) on page 105](#)
- [lacp \(Aggregated Ethernet\) on page 106](#)
- [periodic on page 106](#)

## aggregated-devices

---

<b>Syntax</b>	<pre>aggregated-devices {     ethernet {         device-count <i>number</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit <a href="#">chassis</a> ], [edit <a href="#">chassis</a> node-group <i>name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure properties for aggregated devices on the switch.  The remaining statements are explained separately.
<b>Default</b>	Aggregated devices are disabled.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 29</a></li><li>• <a href="#">Configuring Link Aggregation on page 33</a></li><li>• <i>Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## aggregated-ether-options

**Syntax** The **fcoe-lag** and **mc-ae** statements are not supported on OCX Series switches.

```
aggregated-ether-options {
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  ethernet-switch-profile {
    tag-protocol-id;
    (fcoe-lag | no-fcoe-lag);
    (flow-control | no-flow-control);
    lacp mode {
      admin-key key;
      periodic interval;
      system-id mac-address;
      force-up;
    }
  }
  (link-protection | no-link-protection);
  link-speed speed;
  local-bias (edit interfaces ae);
  (loopback | no-loopback);
  mc-ae {
    chassis-id chassis-id;
    mc-ae-id mc-ae-id;
    mode (active-active);
    status-control (active | standby);
  }
  minimum-links number;
  rebalance-periodic;
  resilient-hash;
  source-address-filter filter;
  (source-filtering | no-source-filtering);
}
```

**Hierarchy Level** [edit **interfaces** aex]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
Statements **fcoe-lag** and **no-fcoe-lag** introduced in Junos OS Release 13.2X52-D10 for the QFX Series.  
Statements **force-up**, **lacp**, and **resilient-hash** introduced in Junos OS Release 14.1X53-D10 for the QFX Series.  
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Configure properties specific to a specific aggregated Ethernet interface.



**NOTE:** The **fcoe-lag** and **mc-ae** statements are not supported on OCX Series switches.

The statements are explained separately.

**Default** Options are not enabled.

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

**Related Documentation**

- [Understanding Aggregated Ethernet Interfaces and LACP on page 29](#)
- [Configuring Aggregated Ethernet LACP on page 32](#)
- *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*
- *Junos OS Network Interfaces Library for Routing Devices*

## chassis

```
Syntax  chassis {
        routing-engine {
        redundancy {
            failover {
                on-disk-failure {
                disk-failure-action (halt | reboot);
                }
                on-loss-of-keepalives;
            }
            graceful-switchover;
        }
        aggregated-devices {
            ethernet {
                device-count number;
            }
            alarm {
                interface-type {
                alarm-name (red | yellow | ignore);
                }
            }
        }
        forwarding-options profile-name {
            num-65-127-prefix value
        }
        fpc slot {
            auto-speed-detection disable
            pic pic-number{
                port port-number{
                    tunnel-port port-number tunnel-services;
                    channel-speed speed;
                }
                port-range port-range-low port-range-high {
                    channel-speed speed;
                }
            }
        }
        maximum-ecmp next-hops;
    }
```

**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 chassis-specific properties for the switch.  
  
The remaining statements are explained separately.

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

- Related Documentation**
- [Configuring Link Aggregation on page 33](#)

---

## 802.3ad

---

**Syntax**

```
802.3ad aex;  
    lacp {  
        force-up;  
        (primary | backup);  
    }  
    port-priority;  
}
```

**Hierarchy Level** [edit [interfaces](#) *interface-name* [ether-options](#)]

**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** Specify the aggregated Ethernet logical interface number.



**NOTE:** The port-priority statement is not supported on QFabric systems.

---

**Options** **aex**—Aggregated Ethernet logical interface number.

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

- Related Documentation**
- [Configuring Link Aggregation on page 33](#)
  - [Configuring Aggregated Ethernet LACP on page 32](#)
  - [Understanding Aggregated Ethernet Interfaces and LACP on page 29](#)
  - [Troubleshooting an Aggregated Ethernet Interface on page 37](#)
  - *Junos OS Network Interfaces Library for Routing Devices*

## device-count

---

<b>Syntax</b>	<code>device-count <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit <a href="#">chassis aggregated-devices ethernet</a> ], [edit <a href="#">chassis node-group <i>name</i> aggregated-devices ethernet</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the number of aggregated Ethernet logical devices available to the switch.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 33</a></li> <li>• <i>Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch</i></li> </ul>

## ethernet

---

<b>Syntax</b>	<code>ethernet {     <a href="#">device-count <i>number</i></a>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">chassis aggregated-devices</a> ], [edit <a href="#">chassis node-group aggregated-devices</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure properties for aggregated Ethernet devices on the switch.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 33</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>

## force-up

---

<b>Syntax</b>	force-up;
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">ether-options</a> 802.3ad lacp; [edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">aggregated-ether-options</a> lacp;
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the state of the interface as up when the peer has limited LACP capability. You can also configure the peer interface (in MC-LAG) to remain up even with limited LACP capability.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 29</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 32</a></li><li>• <i>Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch</i></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li><li>• </li></ul>



## lacp (802.3ad)

**Syntax** `lacp {  
    force-up;  
    (primary | backup);  
    port-priority;  
}`

**Hierarchy Level** [edit [interfaces interface-name ether-options 802.3ad](#)]

**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 the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.



**NOTE:** The port-priority statement is not supported on QFabric systems.

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

**Related Documentation**

- [Configuring Link Aggregation on page 33](#)
- [Configuring Aggregated Ethernet LACP on page 32](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 29](#)

## larp (Aggregated Ethernet)

---

<b>Syntax</b>	<code>larp (active   passive) {   admin-key <i>key</i>;  <i>periodic interval</i>   system-ID <i>mac-address</i>;   force-up; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces interface-name</a> <a href="#">aggregated-ether-options</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Link Aggregation on page 33</a></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 32</a></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 29</a></li></ul>

## periodic

---

<b>Syntax</b>	<code>periodic (fast   slow);</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces aex</a> <a href="#">aggregated-ether-options larp</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure the interval for periodic transmission of LACP packets.
<b>Default</b>	<b>fast</b>
<b>Options</b>	<i>interval</i> —Interval at which to periodically transmit LACP packets: <ul style="list-style-type: none"><li>• <b>fast</b>—Receive packets every second. This is the default.</li><li>• <b>slow</b>—Receive packets every 30 seconds.</li></ul>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 29</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## CHAPTER 8

# Configuration Statements for IRB Interfaces

- `irb` (Interfaces) on page 108
- `l3-interface` (VLAN) on page 111
- `vlan` (OCX Series) on page 112

## irb (Interfaces)

---

```
Syntax  irb {
    accounting-profile name;
    arp-l2-validate;
    description text;

    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;

    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        encapsulation type;
        family inet {
            accounting {
                destination-class-usage;
                source-class-usage {
                    input;
                    output;
                }
            }
        }
        address ipv4-address {
            arp ip-address (mac | multicast-mac) mac-address <publish>;
            broadcast address;
            preferred;
            primary;
            vrrp-group group-number {
                (accept-data | no-accept-data);
                advertise-interval seconds;
                advertisements-threshold number;
                authentication-key key;
                authentication-type authentication;
                fast-interval milliseconds;
                (preempt | no-preempt) {
                    hold-time seconds;
                }
                priority number;
                track {
                    interface interface-name {
                        bandwidth-threshold bandwidth;
                        priority-cost number;
                    }
                    priority-hold-time seconds;
                    route ip-address/mask routing-instance instance-name priority-cost cost;
                }
            }
        }
    }
}
```

```

    virtual-address [ addresses ];
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
no-neighbor-learn;
no-redirects;
primary;
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
}
family inet6 {
    accounting {
        destination-class-usage;
        source-class-usage {
            input;
            output;
        }
    }
}
address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        advertisements-threshold number;
        authentication-key key;
        authentication-type authentication;
        fast-interval milliseconds;
        inet6-advertise-interval milliseconds;
        preempt | no-preempt {
            hold-time seconds;
        }
        priority number;
        track {
            interface interface-name {
                bandwidth-threshold bandwidth priority-cost number;
                priority-cost number;
            }
            priority-hold-time seconds;
        }
    }
}

```

```

        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
    }
}
}
(dad-disable | no-dad-disable);
filter {
    input filter-name;
    output filter-name;
}
mtu bytes;
nd6-stale-time seconds;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
}
family iso {
    address interface-address;
    mtu bytes;
}
family mpls {
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    policer {
        input policer-name;
        output policer-name;
    }
}
native-inner-vlan-id vlan-id;
proxy-arp (restricted | unrestricted);
(traps | no-traps);
vlan-id-list [vlan-id's];
vlan-id-range [vlan-id-range];
}
}

```

Hierarchy Level [edit interfaces *interface-name*

<b>Release Information</b>	Statement introduced in Junos OS Release 12.3R2 for EX Series switches. <b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Configure the properties of a specific integrated bridging and routing (IRB) interface.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>[edit interfaces] Hierarchy Level</i></li> <li>• <i>[edit interfaces] Configuration Statement Hierarchy on EX Series Switches</i></li> </ul>

## l3-interface (VLAN)

<b>Syntax</b>	<code>l3-interface (vlan.logical-interface-number   irb.logical-interface-number);</code>
<b>Hierarchy Level</b>	<code>[edit vlans vlan-name]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1 for the QFX Series. <b>irb</b> option introduced in Junos OS Release 13.2 for the QFX Series.
<b>Description</b>	Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between VLANs. Traffic between VLANs must be routed, which requires a common Layer 3 interface.
<b>Default</b>	No Layer 3 (routing) interface is associated with the VLAN.
<b>Options</b>	<code>vlan.logical-interface-number</code> —Number of the logical interface. Use the <b>unit</b> number that you used when you created the <b>vlan</b> interface with a <b>set interfaces vlan unit</b> statement.



**NOTE:** Use this statement with versions of Junos OS that do not support Enhanced Layer 2 Software (ELS).

`irb.logical-interface-number`—Logical interface defined with a **set interfaces irb** statement.



**NOTE:** Use this statement with versions of Junos OS that support Enhanced Layer 2 Software (ELS).

<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>show ethernet-switching interfaces</i></li> <li>• <i>show vlans</i></li> </ul>

## vlan (OCX Series)

---

<b>Syntax</b>	<pre>vlan {     vlan-name {         description <i>text-description</i>;         l3-interface irb.<i>logical-unit-number</i>;         vlan-id <i>number</i>;         vlan-id-list [<i>vlan-id</i>   <i>vlan-id-vlan-id</i>];     } }</pre>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced in Junos OS Release 14.1X53-D30 for the OCX Series.
<b>Description</b>	Configure VLAN properties on the OCX Series.
<b>Default</b>	If you use the default factory configuration, all switch interfaces become part of the VLAN default.
<b>Options</b>	<p><b><i>vlan-name</i></b>—Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long.</p> <p>The remaining statements are described separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Understanding Integrated Routing and Bridging on page 41</a></li><li>• <a href="#">Configuring VLANs on page 43</a></li></ul>



## CHAPTER 9

# Interfaces Operational Commands

- `monitor interface`
- `show interfaces diagnostics optics`
- `show interfaces ge`
- `show interfaces (GRE)`
- `show interfaces queue`
- `show interfaces xe`

## monitor interface

**Syntax** `monitor interface`  
`<interface-name> | traffic <detail>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Command introduced in Junos OS Release 9.0 for EX Series switches.  
 Command introduced in Junos OS Release 11.1 for the QFX Series.  
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.



**NOTE:** This command is not supported on the QFX3000 QFabric system.

**Options** **none**—Display real-time statistics for all interfaces.

**detail**—(Optional) With traffic option only, display detailed output.

**interface-name**—(Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

**traffic**—(Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

**Additional Information** The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the c key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type in the [CLI Explorer](#). To control the output of the **monitor interface** command while it is running, use the keys listed in [Table 12 on page 114](#). The keys are not case-sensitive.

**Table 12: Output Control Keys for the monitor interface Command**

Key	Action
c	Clears (returns to zero) the delta counters since <b>monitor interface</b> was started. This does not clear the accumulative counter. To clear the accumulative counter, use the <b>clear interfaces interval</b> command.
f	Freezes the display, halting the display of updated statistics and delta counters.
i	Displays information about a different interface. The command prompts you for the name of a specific interface.

**Table 12: Output Control Keys for the monitor interface Command** (*continued*)

Key	Action
n	Displays information about the next interface. The <b>monitor interface</b> command displays the physical or logical interfaces in the same order as the <b>show interfaces terse</b> command.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 13 on page 115](#). The keys are not case-sensitive.

**Table 13: Output Control Keys for the monitor interface traffic Command**

Key	Action
b	Displays the statistics in units of bits and bits per second (bps).
c	Clears (return to 0) the delta counters in the <b>Current Delta</b> column. The statistics counters are not cleared.
d	Displays the <b>Current Delta</b> column (instead of the rate column) in Bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or Esc	Quits the command and returns to the command prompt.
r	Displays the rate column (instead of the <b>Current Delta</b> column) in Bps and pps.

**Required Privilege Level** trace

**List of Sample Output** [monitor interface \(Physical\) on page 117](#)  
[monitor interface \(OTN Interface\) on page 118](#)  
[monitor interface \(MX2020 Routers with MPC6E and OTN MICInterface\) on page 119](#)  
[monitor interface \(Logical\) on page 120](#)  
[monitor interface \(QFX3500 Switch\) on page 120](#)  
[monitor interface traffic on page 121](#)  
[monitor interface traffic \(QFX3500 Switch\) on page 121](#)  
[monitor interface traffic detail \(QFX3500 Switch\) on page 122](#)

**Output Fields** [Table 14 on page 116](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 14: monitor interface Output Fields

Field Name	Field Description	Level of Output
<b>routerl</b>	Hostname of the router.	All levels
<b>Seconds</b>	How long the monitor interface command has been running or how long since you last cleared the counters.	All levels
<b>Time</b>	Current time (UTC).	All levels
<b>Delay x/y/z</b>	Time difference between when the statistics were displayed and the actual clock time. <ul style="list-style-type: none"> <li>• <b>x</b>—Time taken for the last polling (in milliseconds).</li> <li>• <b>y</b>—Minimum time taken across all pollings (in milliseconds).</li> <li>• <b>z</b>—Maximum time taken across all pollings (in milliseconds).</li> </ul>	All levels
<b>Interface</b>	Short description of the interface, including its name, status, and encapsulation.	All levels
<b>Link</b>	State of the link: <b>Up</b> , <b>Down</b> , or <b>Test</b> .	All levels
<b>Current delta</b>	Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.	All levels
<b>Local Statistics</b>	(Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	All levels
<b>Remote Statistics</b>	(Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	All levels

Table 14: monitor interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	All levels
Description	With the <b>traffic</b> option, displays the interface description configured at the <b>[edit interfaces <i>interface-name</i>]</b> hierarchy level.	detail

## Sample Output

### monitor interface (Physical)

```

user@host> monitor interface so-0/0/0
router1                               Seconds: 19                      Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C48
Traffic statistics:
    Input packets:                6045 (0 pps)
    Input bytes:                  6290065 (0 bps)
    Output packets:               10376 (0 pps)
    Output bytes:                 10365540 (0 bps)
Encapsulation statistics:
    Input keepalives:             1901
    Output keepalives:            1901
    NCP state: Opened
    LCP state: Opened
Error statistics:
    Input errors:                 0
    Input drops:                  0
    Input framing errors:         0
    Policed discards:             0
    L3 incompletes:               0
    L2 channel errors:            0
    L2 mismatch timeouts:         0
    Carrier transitions:          1
    Output errors:                0
    Output drops:                 0
    Aged packets:                 0
Active alarms : None
Active defects: None
SONET error counts/seconds:
    LOS count                     1
    LOF count                     1
    SEF count                     1
    ES-S                          0
    SES-S                         0
SONET statistics:
    BIP-B1                       458871

```

```

BIP-B2                      460072          [0]
REI-L                      465610          [0]
BIP-B3                      458978          [0]
REI-P                      458773          [0]

```

## Received SONET overhead:

```

F1      : 0x00 J0      : 0x00 K1      : 0x00
K2      : 0x00 S1      : 0x00 C2      : 0x00
C2(cmp) : 0x00 F2      : 0x00 Z3      : 0x00
Z4      : 0x00 S1(cmp) : 0x00

```

## Transmitted SONET overhead:

```

F1      : 0x00 J0      : 0x01 K1      : 0x00
K2      : 0x00 S1      : 0x00 C2      : 0xcf
F2      : 0x00 Z3      : 0x00 Z4      : 0x00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

## monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0
```

```

Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:                0 (0 bps)
  Output bytes:               0 (0 bps)
  Input packets:              0 (0 pps)
  Output packets:             0 (0 pps)
Error statistics:
  Input errors:                0
  Input drops:                 0
  Input framing errors:        0
  Policed discards:            0
  L3 incompletes:              0
  L2 channel errors:           0
  L2 mismatch timeouts:        0
  Carrier transitions:          5
  Output errors:               0
  Output drops:                0
  Aged packets:                0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets              0
  Broadcast packets            0
  Multicast packets            0
  Oversized frames             0
  Packet reject count          0
  DA rejects                   0
  SA rejects                   0
Output MAC/Filter Statistics:
  Unicast packets              0
  Broadcast packets            0
  Multicast packets            0
  Packet pad count             0
  Packet error count           0
OTN Link 0
  OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
  OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
  OTN OC - Seconds
    LOS                        2

```

```

      LOF                                9
OTN OTU - FEC Statistics
  Corr err ratio                        N/A
  Corr bytes                           0
  Uncorr words                          0
OTN OTU - Counters
  BIP                                  0
  BBE                                  0
  ES                                   0
  SES                                  0
  UAS                                  422
OTN ODU - Counters
  BIP                                  0
  BBE                                  0
  ES                                   0
  SES                                  0
  UAS                                  422
OTN ODU - Received Overhead    APSGCC 0-3:      0

```

### monitor interface (MX2020 Routers with MPC6E and OTN MICInterface)

```

user@host> monitor interface xe-3/0/0
host name                               Seconds: 67                                Time: 23:46:46
                                                                 Delay: 0/0/13

Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:                                Current delta
  Input bytes:                                0 (0 bps)                [0]
  Output bytes:                               0 (0 bps)                [0]
  Input packets:                              0 (0 pps)                [0]
  Output packets:                             0 (0 pps)                [0]
Error statistics:
  Input errors:                               0                        [0]
  Input drops:                                0                        [0]
  Input framing errors:                       0                        [0]
  Policed discards:                           0                        [0]
  L3 incompletes:                             0                        [0]
  L2 channel errors:                          0                        [0]
  L2 mismatch timeouts:                       0                        [0]
  Carrier transitions:                         3                        [0]
  Output errors:                              0                        [0]
  Output drops:                               0                        [0]
  Aged packets:                              0                        [0]
OTN Link 0
OTN Alarms:
OTN Defects:
OTN OC - Seconds
  LOS                                         0                        [0]
  LOF                                         0                        [0]
OTN OTU - FEC Statistics
  Corr err ratio                            N/A
  Corr bytes                                0                        [0]
  Uncorr words                              0                        [0]
OTN OTU - Counters
  BIP                                         0                        [0]
  BBE                                         0                        [0]
  ES                                          0                        [0]
  SES                                          0                        [0]
  UAS                                          0                        [0]
OTN ODU - Counters
  BIP                                         0

```

```

BBE                                0                                [0]
ES                                0                                [0]
SES                                0                                [0]
UAS                                0                                [0]
OTN ODU - Received Overhead      [0]
APSPCC 0-3:                      00 00 00 00

```

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

### monitor interface (Logical)

```

user@host> monitor interface so-1/0/0.0
host name                Seconds: 16                Time: 15:33:39
                                                Delay: 0/0/1

Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:
Input bytes:              0                                [0]
Output bytes:             0                                [0]
Input packets:            0                                [0]
Output packets:           0                                [0]
Remote statistics:
Input bytes:              0 (0 bps)                       [0]
Output bytes:             0 (0 bps)                       [0]
Input packets:            0 (0 pps)                       [0]
Output packets:           0 (0 pps)                       [0]
Traffic statistics:
Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'

```

### monitor interface (QFX3500 Switch)

```

user@switch> monitor interface ge-0/0/0
Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:
Input bytes:              0 (0 bps)                       [0]
Output bytes:             0 (0 bps)                       [0]
Input packets:            0 (0 pps)                       [0]
Output packets:           0 (0 pps)                       [0]
Error statistics:
Input errors:             0                                [0]
Input drops:              0                                [0]
Input framing errors:     0                                [0]
Policed discards:        0                                [0]
L3 incompletes:          0                                [0]
L2 channel errors:       0                                [0]
L2 mismatch timeouts:    0                                [0]
Carrier transitions:      0                                [0]
Output errors:            0                                [0]
Output drops:             0                                [0]
Aged packets:            0                                [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
Unicast packets          0                                [0]
Broadcast packets        0 Multicast packet              [0]

```



Interface warnings:  
 o Outstanding LINK alarm

### monitor interface traffic

```
user@host> monitor interface traffic
host name                               Seconds: 15                               Time: 12:31:09
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
so-1/0/0	Down	0	(0)	0	(0)
so-1/1/0	Down	0	(0)	0	(0)
so-1/1/1	Down	0	(0)	0	(0)
so-1/1/2	Down	0	(0)	0	(0)
so-1/1/3	Down	0	(0)	0	(0)
t3-1/2/0	Down	0	(0)	0	(0)
t3-1/2/1	Down	0	(0)	0	(0)
t3-1/2/2	Down	0	(0)	0	(0)
t3-1/2/3	Down	0	(0)	0	(0)
so-2/0/0	Up	211035	(1)	36778	(0)
so-2/0/1	Up	192753	(1)	36782	(0)
so-2/0/2	Up	211020	(1)	36779	(0)
so-2/0/3	Up	211029	(1)	36776	(0)
so-2/1/0	Up	189378	(1)	36349	(0)
so-2/1/1	Down	0	(0)	18747	(0)
so-2/1/2	Down	0	(0)	16078	(0)
so-2/1/3	Up	0	(0)	80338	(0)
at-2/3/0	Up	0	(0)	0	(0)
at-2/3/1	Down	0	(0)	0	(0)

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D

### monitor interface traffic (QFX3500 Switch)

```
user@switch> monitor interface traffic
switch                               Seconds: 7                               Time: 16:04:37
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
ge-0/0/0	Down	0	(0)	0	(0)
ge-0/0/1	Up	392187	(0)	392170	(0)
ge-0/0/2	Down	0	(0)	0	(0)
ge-0/0/3	Down	0	(0)	0	(0)
ge-0/0/4	Down	0	(0)	0	(0)
ge-0/0/5	Down	0	(0)	0	(0)
ge-0/0/6	Down	0	(0)	0	(0)
ge-0/0/7	Down	0	(0)	0	(0)
ge-0/0/8	Down	0	(0)	0	(0)
ge-0/0/9	Up	392184	(0)	392171	(0)
ge-0/0/10	Down	0	(0)	0	(0)
ge-0/0/11	Down	0	(0)	0	(0)
ge-0/0/12	Down	0	(0)	0	(0)
ge-0/0/13	Down	0	(0)	0	(0)
ge-0/0/14	Down	0	(0)	0	(0)
ge-0/0/15	Down	0	(0)	0	(0)
ge-0/0/16	Down	0	(0)	0	(0)
ge-0/0/17	Down	0	(0)	0	(0)
ge-0/0/18	Down	0	(0)	0	(0)
ge-0/0/19	Down	0	(0)	0	(0)
ge-0/0/20	Down	0	(0)	0	(0)
ge-0/0/21	Down	0	(0)	0	(0)
ge-0/0/22	Up	392172	(0)	392187	(0)
ge-0/0/23	Up	392185	(0)	392173	(0)

vcp-0	Down	0		0	
vcp-1	Down	0		0	
ae0	Down	0	(0)	0	(0)
bme0	Up	0		1568706	

### monitor interface traffic detail (QFX3500 Switch)

user@switch> monitor interface traffic detail  
switch

Seconds: 74

Time: 16:03:02

Interface Description	Link	Input packets	(pps)	Output packets	(pps)
ge-0/0/0	Down	0	(0)	0	(0)
ge-0/0/1	Up	392183	(0)	392166	(0)
ge-0/0/2	Down	0	(0)	0	(0)
ge-0/0/3	Down	0	(0)	0	(0)
ge-0/0/4	Down	0	(0)	0	(0)
ge-0/0/5	Down	0	(0)	0	(0)
ge-0/0/6	Down	0	(0)	0	(0)
ge-0/0/7	Down	0	(0)	0	(0)
ge-0/0/8	Down	0	(0)	0	(0)
ge-0/0/9	Up	392181	(0)	392168	(0)
ge-0/0/10	Down	0	(0)	0	(0)
ge-0/0/11	Down	0	(0)	0	(0)
ge-0/0/12	Down	0	(0)	0	(0)
ge-0/0/13	Down	0	(0)	0	(0)
ge-0/0/14	Down	0	(0)	0	(0)
ge-0/0/15	Down	0	(0)	0	(0)
ge-0/0/16	Down	0	(0)	0	(0)
ge-0/0/17	Down	0	(0)	0	(0)
ge-0/0/18	Down	0	(0)	0	(0)
ge-0/0/19	Down	0	(0)	0	(0)
ge-0/0/20	Down	0	(0)	0	(0)
ge-0/0/21	Down	0	(0)	0	(0)
ge-0/0/22	Up	392169	(0)	392184	(1)
ge-0/0/23	Up	392182	(0)	392170	(0)
vcp-0	Down	0		0	
vcp-1	Down	0		0	
ae0	Down	0	(0)	0	(0)
bme0	Up	0		1568693	

## show interfaces diagnostics optics

<b>Syntax</b>	<code>show interfaces diagnostics optics <i>interface-name</i></code>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	<p>Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, XFP, QSFP+, or CFP) installed in EX Series or QFX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information. For a list of transceivers supported on EX Series switches and their specifications, including DOM support, see <i>Pluggable Transceivers Supported on EX Series Switches</i>.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p>
<b>Options</b>	<i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed: <i>ge-fpc/pic/port</i> , <i>xe-fpc/pic/port</i> , or <i>et-fpc/pic/port</i> .
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Monitoring Interface Status and Traffic</i></li> <li>• <a href="#">Monitoring Interface Status and Traffic on page 17</a></li> <li>• <i>Installing a Transceiver in an EX Series Switch</i></li> <li>• <i>Installing a Transceiver in a QFX Series Device</i></li> <li>• <i>Removing a Transceiver from a Switch</i></li> <li>• <i>Removing a Transceiver from a QFX Series Device</i></li> <li>• <a href="#">Junos OS Ethernet Interfaces Configuration Guide</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 130</a></p> <p><a href="#">show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 131</a></p> <p><a href="#">show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 132</a></p> <p><a href="#">show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver) on page 133</a></p> <p><a href="#">show interfaces diagnostics optics et-4/1/0 (CFP Transceiver) on page 134</a></p>
<b>Output Fields</b>	Table 15 on page 123 lists the output fields for the <b>show interfaces diagnostics optics</b> command. Output fields are listed in the approximate order in which they appear.

Table 15: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.

Table 15: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Laser bias current</b>	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
<b>Laser output power</b> (Not available for QSFP+ transceivers)	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
<b>Laser temperature</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the laser temperature, in Celsius and Fahrenheit.
<b>Module temperature</b>	Displays the temperature, in Celsius and Fahrenheit.
<b>Module voltage</b> (Not available for XFP transceivers)	Displays the voltage, in Volts.
<b>Laser rx power</b> (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Displays the laser received optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
<b>Receiver signal average optical power</b> (Not available for XFP, QSFP+, and CFP transceivers)	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
<b>Laser bias current high alarm</b>	Displays whether the laser bias power setting high alarm is <b>On</b> or <b>Off</b> .
<b>Laser bias current low alarm</b>	Displays whether the laser bias power setting low alarm is <b>On</b> or <b>Off</b> .
<b>Laser bias current high warning</b>	Displays whether the laser bias power setting high warning is <b>On</b> or <b>Off</b> .
<b>Laser bias current low warning</b>	Displays whether the laser bias power setting low warning is <b>On</b> or <b>Off</b> .
<b>Laser output power high alarm</b> (Not available for QSFP+ transceivers)	Displays whether the laser output power high alarm is <b>On</b> or <b>Off</b> .
<b>Laser output power low alarm</b> (Not available for QSFP+ transceivers)	Displays whether the laser output power low alarm is <b>On</b> or <b>Off</b> .
<b>Laser output power high warning</b> (Not available for QSFP+ transceivers)	Displays whether the laser output power high warning is <b>On</b> or <b>Off</b> .

Table 15: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Laser output power low warning</b> (Not available for QSFP+ transceivers)	Displays whether the laser output power low warning is <b>On</b> or <b>Off</b> .
<b>Laser temperature high alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature high alarm is <b>On</b> or <b>Off</b> .
<b>Laser temperature low alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature low alarm is <b>On</b> or <b>Off</b> .
<b>Laser temperature high warning</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature high warning is <b>On</b> or <b>Off</b> .
<b>Laser temperature low warning</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the laser temperature low warning is <b>On</b> or <b>Off</b> .
<b>Module temperature high alarm</b> (Not available for QSFP+ transceivers)	Displays whether the module temperature high alarm is <b>On</b> or <b>Off</b> .
<b>Module temperature low alarm</b> (Not available for QSFP+ transceivers)	Displays whether the module temperature low alarm is <b>On</b> or <b>Off</b> .
<b>Module temperature high warning</b> (Not available for QSFP+ transceivers)	Displays whether the module temperature high warning is <b>On</b> or <b>Off</b> .
<b>Module temperature low warning</b> (Not available for QSFP+ transceivers)	Displays whether the module temperature low warning is <b>On</b> or <b>Off</b> .
<b>Module voltage high alarm</b> (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage high alarm is <b>On</b> or <b>Off</b> .
<b>Module voltage low alarm</b> (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage low alarm is <b>On</b> or <b>Off</b> .
<b>Module voltage high warning</b> (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage high warning is <b>On</b> or <b>Off</b> .
<b>Module voltage low warning</b> (Not available for XFP and QSFP+ transceivers)	Displays whether the module voltage low warning is <b>On</b> or <b>Off</b> .

Table 15: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Laser rx power high alarm</b> (Not available for QSFP+ and CFP transceivers)	Displays whether the receive laser power high alarm is <b>On</b> or <b>Off</b> .
<b>Laser rx power low alarm</b> (Not available for QSFP+ and CFP transceivers)	Displays whether the receive laser power low alarm is <b>On</b> or <b>Off</b> .
<b>Laser rx power high warning</b> (Not available for QSFP+ and CFP transceivers)	Displays whether the receive laser power high warning is <b>On</b> or <b>Off</b> .
<b>Laser rx power low warning</b> (Not available for QSFP+ and CFP transceivers)	Displays whether the receive laser power low warning is <b>On</b> or <b>Off</b> .
<b>Laser bias current high alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser bias current high alarm.
<b>Module not ready alarm</b> (Not available for SFP, SFP+, and QSFP+ transceivers)	Displays whether the module not ready alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , the module has an operational fault.
<b>Module low power alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the module low power alarm is <b>On</b> or <b>Off</b> .
<b>Module initialization incomplete alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the module initialization incomplete alarm is <b>On</b> or <b>Off</b> .
<b>Module fault alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the module fault alarm is <b>On</b> or <b>Off</b> .
<b>PLD Flash initialization fault alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the PLD Flash initialization fault alarm is <b>On</b> or <b>Off</b> .
<b>Power supply fault alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the power supply fault alarm is <b>On</b> or <b>Off</b> .
<b>Checksum fault alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the checksum fault alarm is <b>On</b> or <b>Off</b> .
<b>Tx laser disabled alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the Tx laser disabled alarm is <b>On</b> or <b>Off</b> .

Table 15: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Module power down alarm</b> (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Displays whether the module power down alarm is <b>On</b> or <b>Off</b> . When the output is <b>On</b> , module is in a limited power mode, low for normal operation.
<b>Tx data not ready alarm</b> (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx data not ready alarm is <b>On</b> or <b>Off</b> .
<b>Tx not ready alarm</b> (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Any condition leading to invalid data on the transmit path. Displays whether the Tx not ready alarm is <b>On</b> or <b>Off</b> .
<b>Tx laser fault alarm</b> (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Laser fault condition. Displays whether the Tx laser fault alarm is <b>On</b> or <b>Off</b> .
<b>Tx CDR loss of lock alarm</b> (Not available for SFP, SFP+, and QSFP+ transceivers)	Transmit clock and data recovery (CDR) loss of lock. Loss of lock on the transmit side of the CDR. Displays whether the Tx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .
<b>Rx not ready alarm</b> (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Any condition leading to invalid data on the receive path. Displays whether the Rx not ready alarm is <b>On</b> or <b>Off</b> .
<b>Rx loss of signal alarm</b> (Not available for SFP and SFP+ transceivers)	Receive loss of signal alarm. When the output is <b>On</b> , indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is <b>On</b> or <b>Off</b> .
<b>Rx CDR loss of lock alarm</b> (Not available for SFP, SFP+, and QSFP+ transceivers)	Receive CDR loss of lock. Loss of lock on the receive side of the CDR. Displays whether the Rx CDR loss of lock alarm is <b>On</b> or <b>Off</b> .
<b>Laser bias current low alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser bias current low alarm.
<b>Laser bias current high warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser bias current high warning.
<b>Laser bias current low warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser bias current low warning.
<b>Laser output power high alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser output power high alarm.
<b>Laser output power low alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser output power low alarm.

Table 15: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Laser output power high warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser output power high warning.
<b>Laser output power low warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser output power low warning.
<b>Module temperature high alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the module temperature high alarm.
<b>Module temperature low alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the module temperature low alarm.
<b>Module temperature high warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the module temperature high warning.
<b>Module temperature low warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the module temperature low warning.
<b>Module voltage high alarm threshold</b> (Not available for XFP and QSFP+ transceivers)	Displays the vendor-specified threshold for the module voltage high alarm.
<b>Module voltage low alarm threshold</b> (Not available for XFP and QSFP+ transceivers)	Displays the vendor-specified threshold for the module voltage low alarm.
<b>Module voltage high warning threshold</b> (Not available for XFP and QSFP+ transceivers)	Displays the vendor-specified threshold for the module voltage high warning.
<b>Module voltage low warning threshold</b> (Not available for XFP and QSFP+ transceivers)	Displays the vendor-specified threshold for the module voltage low warning.
<b>Laser rx power high alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser rx power high alarm.
<b>Laser rx power low alarm threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser rx power low alarm.
<b>Laser rx power high warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser rx power high warning.



Table 15: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Laser rx power low warning threshold</b> (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser rx power low warning.
<b>Laser temperature high alarm threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for the laser temperature high alarm, in Celsius and Fahrenheit.
<b>Laser temperature low alarm threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for the laser temperature low alarm, in Celsius and Fahrenheit.
<b>Laser temperature high warning threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for the laser temperature high warning, in Celsius and Fahrenheit.
<b>Laser temperature low warning threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for the laser temperature low warning, in Celsius and Fahrenheit.
<b>SOA bias current high alarm threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for SOA bias current high alarm.
<b>SOA bias current low alarm threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for SOA bias current low alarm.
<b>SOA bias current high warning threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for SOA bias current high warning.
<b>SOA bias current low warning threshold</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays the vendor-specified threshold for SOA bias current low warning.
<b>Laser receiver power high alarm</b> (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power high alarm is <b>On</b> or <b>Off</b> .
<b>Laser receiver power low alarm</b> (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power low alarm is <b>On</b> or <b>Off</b> .
<b>Laser receiver power high warning</b> (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power high warning is <b>On</b> or <b>Off</b> .
<b>Laser receiver power low warning</b> (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the laser receiver power low warning is <b>On</b> or <b>Off</b> .

Table 15: show interfaces diagnostics optics Output Fields (*continued*)

Field Name	Field Description
<b>Laser receiver power</b> (Not available for SFP, SFP+, and XFP transceivers)	Displays the laser receiver power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
<b>Tx loss of signal functionality alarm</b> (Not available for SFP, SFP+, and XFP transceivers)	Displays whether the Tx loss of signal functionality alarm is <b>On</b> or <b>Off</b> .
<b>APD supply fault alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the APD supply fault alarm is <b>On</b> or <b>Off</b> .
<b>TEC fault alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the TEC fault alarm is <b>On</b> or <b>Off</b> .
<b>Wavelength unlocked alarm</b> (Not available for SFP, SFP+, XFP, and QSFP+ transceivers)	Displays whether the Wavelength unlocked alarm is <b>On</b> or <b>Off</b> .

## Sample Output

### show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver)

```

user@switch> show interfaces diagnostics optics ge-0/1/0
Physical interface: ge-0/1/0
  Laser bias current           : 5.444 mA
  Laser output power          : 0.3130 mW / -5.04 dBm
  Module temperature          : 36 degrees C / 97 degrees F
  Module voltage              : 3.2120 V
  Receiver signal average optical power : 0.3840 mW / -4.16 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser output power high alarm  : Off
  Laser output power low alarm   : Off
  Laser output power high warning : Off
  Laser output power low warning : Off
  Module temperature high alarm  : Off
  Module temperature low alarm   : Off
  Module temperature high warning : Off
  Module temperature low warning : Off
  Module voltage high alarm      : Off
  Module voltage low alarm       : Off
  Module voltage high warning    : Off
  Module voltage low warning     : Off
  Laser rx power high alarm      : Off
  Laser rx power low alarm       : Off
  Laser rx power high warning    : Off
  Laser rx power low warning     : Off
  Laser bias current high alarm threshold : 15.000 mA
  Laser bias current low alarm threshold  : 1.000 mA
  Laser bias current high warning threshold : 12.000 mA

```

```

Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6300 mW / -2.01 dBm
Laser output power low alarm threshold : 0.0660 mW / -11.80 dBm
Laser output power high warning threshold : 0.6300 mW / -2.01 dBm
Laser output power low warning threshold : 0.0780 mW / -11.08 dBm
Module temperature high alarm threshold : 109 degrees C / 228 degrees F
Module temperature low alarm threshold : -29 degrees C / -20 degrees F
Module temperature high warning threshold : 103 degrees C / 217 degrees F
Module temperature low warning threshold : -13 degrees C / 9 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2589 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7939 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0157 mW / -18.04 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
Laser bias current : 4.968 mA
Laser output power : 0.4940 mW / -3.06 dBm
Module temperature : 27 degrees C / 81 degrees F
Module voltage : 3.2310 V
Receiver signal average optical power : 0.0000
Laser bias current high alarm : Off
Laser bias current low alarm : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm : Off
Module voltage low alarm : Off
Module voltage high warning : Off
Module voltage low warning : Off
Laser rx power high alarm : Off
Laser rx power low alarm : On
Laser rx power high warning : Off
Laser rx power low warning : On
Laser bias current high alarm threshold : 10.500 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 9.000 mA
Laser bias current low warning threshold : 2.500 mA
Laser output power high alarm threshold : 1.4120 mW / 1.50 dBm
Laser output power low alarm threshold : 0.0740 mW / -11.31 dBm
Laser output power high warning threshold : 0.7070 mW / -1.51 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 75 degrees C / 167 degrees F
Module temperature low alarm threshold : -5 degrees C / 23 degrees F
Module temperature high warning threshold : 70 degrees C / 158 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F

```

```

Module voltage high alarm threshold      : 3.630 V
Module voltage low alarm threshold       : 2.970 V
Module voltage high warning threshold    : 3.465 V
Module voltage low warning threshold     : 3.135 V
Laser rx power high alarm threshold      : 1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold       : 0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver)

```

user@switch> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
Laser bias current                : 8.029 mA
Laser output power                 : 0.6430 mW / -1.92 dBm
Module temperature                 : 4 degrees C / 39 degrees F
Laser rx power                     : 0.0012 mW / -29.21 dBm
Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Module temperature high alarm      : Off
Module temperature low alarm       : Off
Module temperature high warning    : Off
Module temperature low warning     : Off
Laser rx power high alarm          : Off
Laser rx power low alarm           : On
Laser rx power high warning        : Off
Laser rx power low warning         : On
Module not ready alarm             : On
Module power down alarm            : Off
Tx data not ready alarm            : Off
Tx not ready alarm                 : Off
Tx laser fault alarm              : Off
Tx CDR loss of lock alarm          : Off
Rx not ready alarm                 : On
Rx loss of signal alarm            : On
Rx CDR loss of lock alarm          : On
Laser bias current high alarm threshold : 13.000 mA
Laser bias current low alarm threshold : 2.000 mA
Laser bias current high warning threshold : 12.000 mA
Laser bias current low warning threshold : 3.000 mA
Laser output power high alarm threshold : 0.8310 mW / -0.80 dBm
Laser output power low alarm threshold : 0.1650 mW / -7.83 dBm
Laser output power high warning threshold : 0.7410 mW / -1.30 dBm
Laser output power low warning threshold : 0.1860 mW / -7.30 dBm
Module temperature high alarm threshold : 90 degrees C / 194 degrees F
Module temperature low alarm threshold : 0 degrees C / 32 degrees F
Module temperature high warning threshold : 85 degrees C / 185 degrees F
Module temperature low warning threshold : 0 degrees C / 32 degrees F
Laser rx power high alarm threshold : 0.8912 mW / -0.50 dBm
Laser rx power low alarm threshold : 0.0912 mW / -10.40 dBm
Laser rx power high warning threshold : 0.7943 mW / -1.00 dBm
Laser rx power low warning threshold : 0.1023 mW / -9.90 dBm

```

## Sample Output

### show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver)

```

user@switch> show interfaces diagnostics optics et-3/0/0
Physical interface: et-3/0/0
  Module temperature           : 33 degrees C / 92 degrees F
  Module voltage               : 3.3060 V
Lane 0
  Laser bias current           : 7.182 mA
  Laser receiver power         : 0.743 mW / -1.29 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm         : Off
Lane 1
  Laser bias current           : 7.326 mA
  Laser receiver power         : 0.752 mW / -1.24 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm         : Off
Lane 2
  Laser bias current           : 7.447 mA
  Laser receiver power         : 0.790 mW / -1.03 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm         : Off
Lane 3
  Laser bias current           : 7.734 mA
  Laser receiver power         : 0.768 mW / -1.15 dBm
  Laser bias current high alarm : Off
  Laser bias current low alarm  : Off
  Laser bias current high warning : Off
  Laser bias current low warning : Off
  Laser receiver power high alarm : Off
  Laser receiver power low alarm  : Off
  Laser receiver power high warning : Off
  Laser receiver power low warning : Off
  Tx loss of signal functionality alarm : Off
  Rx loss of signal alarm         : Off

```

## Sample Output

### show interfaces diagnostics optics et-4/1/0 (CFP Transceiver)

```

user@switch> show interfaces diagnostics optics et-4/1/0
Physical interface: et-4/1/0
  Module temperature                : 38 degrees C / 101 degrees F
  Module voltage                    : 3.2500 V
  Module temperature high alarm     : Off
  Module temperature low alarm      : Off
  Module temperature high warning   : Off
  Module temperature low warning    : Off
  Module voltage high alarm         : Off
  Module voltage low alarm          : Off
  Module voltage high warning       : Off
  Module voltage low warning        : Off
  Module not ready alarm            : Off
  Module low power alarm            : Off
  Module initialization incomplete alarm : Off
  Module fault alarm                : Off
  PLD Flash initialization fault alarm : Off
  Power supply fault alarm          : Off
  Checksum fault alarm              : Off
  Tx laser disabled alarm           : Off
  Tx loss of signal functionality alarm : Off
  Tx CDR loss of lock alarm         : Off
  Rx loss of signal alarm           : Off
  Rx CDR loss of lock alarm         : Off
  Module temperature high alarm threshold : 75 degrees C / 167 degrees F
  Module temperature low alarm threshold : -5 degrees C / 23 degrees F
  Module temperature high warning threshold : 70 degrees C / 158 degrees F
  Module temperature low warning threshold : 0 degrees C / 32 degrees F
  Module voltage high alarm threshold : 3.5000 V
  Module voltage low alarm threshold : 3.0990 V
  Module voltage high warning threshold : 3.4000 V
  Module voltage low warning threshold : 3.2000 V
  Laser bias current high alarm threshold : 250.000 mA
  Laser bias current low alarm threshold : 37.500 mA
  Laser bias current high warning threshold : 225.000 mA
  Laser bias current low warning threshold : 50.000 mA
  Laser output power high alarm threshold : 3.9800 mW / 6.00 dBm
  Laser output power low alarm threshold : 0.4670 mW / -3.31 dBm
  Laser output power high warning threshold : 3.5480 mW / 5.50 dBm
  Laser output power low warning threshold : 0.5240 mW / -2.81 dBm
  Laser rx power high alarm threshold : 3.5481 mW / 5.50 dBm
  Laser rx power low alarm threshold : 0.0616 mW / -12.10 dBm
  Laser rx power high warning threshold : 3.1622 mW / 5.00 dBm
  Laser rx power low warning threshold : 0.0691 mW / -11.61 dBm
  Laser temperature high alarm threshold : 67 degrees C / 153 degrees F
  Laser temperature low alarm threshold : 35 degrees C / 95 degrees F
  Laser temperature high warning threshold : 62 degrees C / 144 degrees F
  Laser temperature low warning threshold : 40 degrees C / 104 degrees F
  SOA bias current high alarm threshold : 0.000 mA
  SOA bias current low alarm threshold : 0.000 mA
  SOA bias current high warning threshold : 0.000 mA
  SOA bias current low warning threshold : 0.000 mA
Lane 0
  Laser bias current                : 131.684 mA
  Laser output power                 : 1.002 mW / 0.01 dBm
  Laser temperature                  : 54 degrees C / 128 degrees F
  Laser receiver power               : 0.497 mW / -3.03 dBm

```

```

Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Laser temperature high alarm       : Off
Laser temperature low alarm        : Off
Laser temperature high warning     : Off
Laser temperature low warning      : Off
Laser receiver power high alarm    : Off
Laser receiver power low alarm     : Off
Laser receiver power high warning  : Off
Laser receiver power low warning   : Off
Tx loss of signal functionality alarm : Off
Rx CDR loss of lock alarm          : Off
Rx loss of signal alarm            : Off
Rx CDR loss of lock alarm          : Off
APD supply fault alarm             : Off
TEC fault alarm                   : Off
Wavelength unlocked alarm          : Off

Lane 1
Laser bias current                 : 122.345 mA
Laser output power                 : 1.002 mW / 0.01 dBm
Laser temperature                  : 51 degrees C / 124 degrees F
Laser receiver power               : 0.611 mW / -2.14 dBm
Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off
Laser output power low alarm       : Off
Laser output power high warning    : Off
Laser output power low warning     : Off
Laser temperature high alarm       : Off
Laser temperature low alarm        : Off
Laser temperature high warning     : Off
Laser temperature low warning      : Off
Laser receiver power high alarm    : Off
Laser receiver power low alarm     : Off
Laser receiver power high warning  : Off
Laser receiver power low warning   : Off
Tx loss of signal functionality alarm : Off
Tx CDR loss of lock alarm          : Off
Rx loss of signal alarm            : Off
Rx CDR loss of lock alarm          : Off
APD supply fault alarm             : Off
TEC fault alarm                   : Off
Wavelength unlocked alarm          : Off

Lane 2
Laser bias current                 : 112.819 mA
Laser output power                 : 1.000 mW / 0.00 dBm
Laser temperature                  : 50 degrees C / 122 degrees F
Laser receiver power               : 0.540 mW / -2.67 dBm
Laser bias current high alarm      : Off
Laser bias current low alarm       : Off
Laser bias current high warning    : Off
Laser bias current low warning     : Off
Laser output power high alarm      : Off

```

Laser output power low alarm	: Off
Laser output power high warning	: Off
Laser output power low warning	: Off
Laser temperature high alarm	: Off
Laser temperature low alarm	: Off
Laser temperature high warning	: Off
Laser temperature low warning	: Off
Laser receiver power high alarm	: Off
Laser receiver power low alarm	: Off
Laser receiver power high warning	: Off
Laser receiver power low warning	: Off
Tx loss of signal functionality alarm	: Off
Tx CDR loss of lock alarm	: Off
Rx loss of signal alarm	: Off
Rx CDR loss of lock alarm	: Off
APD supply fault alarm	: Off
TEC fault alarm	: Off
Wavelength unlocked alarm	: Off
Lane 3	
Laser bias current	: 100.735 mA
Laser output power	: 1.002 mW / 0.01 dBm
Laser temperature	: 50 degrees C / 122 degrees F
Laser receiver power	: 0.637 mW / -1.96 dBm
Laser bias current high alarm	: Off
Laser bias current low alarm	: Off
Laser bias current high warning	: Off
Laser bias current low warning	: Off
Laser output power high alarm	: Off
Laser output power low alarm	: Off
Laser output power high warning	: Off
Laser output power low warning	: Off
Laser temperature high alarm	: Off
Laser temperature low alarm	: Off
Laser temperature high warning	: Off
Laser temperature low warning	: Off
Laser receiver power high alarm	: Off
Laser receiver power low alarm	: Off
Laser receiver power high warning	: Off
Laser receiver power low warning	: Off
Tx loss of signal functionality alarm	: Off
Tx CDR loss of lock alarm	: Off
Rx loss of signal alarm	: Off
Rx CDR loss of lock alarm	: Off
APD supply fault alarm	: Off
TEC fault alarm	: Off
Wavelength unlocked alarm	: Off



## show interfaces ge

<b>Syntax</b>	<pre>show interfaces <i>device-name:type-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance (all   <i>instance-name</i>)&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Display status information about the specified Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.
<b>Options</b>	<p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b><i>device-name:type-fpc/pic/port</i></b>—The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing-instance or display all routing-instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 17</a></li> <li>• <a href="#">Troubleshooting Network Interfaces on page 17</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 37</a></li> <li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces on page 145</a></p> <p><a href="#">show interfaces brief on page 145</a></p> <p><a href="#">show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled) on page 145</a></p> <p><a href="#">show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled) on page 146</a></p>

[show interfaces extensive \(Symmetric Flow Control and Autonegotiation Enabled\) on page 147](#)

[show interfaces extensive \(Asymmetric Flow Control and Autonegotiation Enabled\) on page 149](#)

[show interfaces terse on page 151](#)

[show interfaces terse \(QFabric Systems\) on page 151](#)

**Output Fields** Table 16 on page 138 lists the output fields for the **show interfaces ge** command. Output fields are listed in the approximate order in which they appear.

**Table 16: show interfaces ge Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Description</b>	Optional user-specified description.	<b>brief detail extensive</b>
<b>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>MTU</b>	Maximum transmission unit size on the physical interface. The default is 1514.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Loopback</b>	Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .	All levels
<b>Source filtering</b>	Source filtering status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Flow control</b>	Flow control status: <b>Enabled</b> or <b>Disabled</b> .  <i>NOTE:</i> This field is only displayed if asymmetric flow control is not configured.	<b>detail extensive</b>
<b>Configured-flow-control</b>	Configured flow control for the interface transmit buffers ( <b>tx-buffers</b> ) and receive buffers ( <b>rx-buffers</b> ):  <ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer.  <b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer.  <b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul> <i>NOTE:</i> This field is only displayed if asymmetric flow control is configured.	<b>detail extensive</b>

Table 16: show interfaces ge Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Auto-negotiation</b>	Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Remote-fault</b>	Remote fault status: <ul style="list-style-type: none"> <li>• <b>Online</b>—Autonegotiation is manually configured as online.</li> <li>• <b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>	All levels
<b>Device flags</b>	Information about the physical device.	All levels
<b>Interface flags</b>	Information about the interface.	All levels
<b>Link flags</b>	Information about the link.	All levels
<b>CoS queues</b>	Number of CoS queues configured.	<b>detail extensive none</b>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive none</b>
<b>Hardware address</b>	MAC address of the hardware.	<b>detail extensive none</b>
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .	<b>detail extensive none</b>
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled on the switch.</p>	<b>detail extensive</b>

Table 16: show interfaces ge Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Input errors</b>	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>

Table 16: show interfaces ge Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output errors</b>	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	Total number of egress queues supported on the specified interface.	<b>detail extensive</b>
<b>Queue counters (Egress )</b>	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>detail extensive</b>
<b>Queue Number</b>	The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.	<b>detail extensive</b>
<b>Active alarms and Active defects</b>	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>	<b>detail extensive none</b>

Table 16: show interfaces ge Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>MAC statistics</b>	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets</b> and <b>total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul>	<b>extensive</b>
<b>Filter Statistics</b>	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.	<b>extensive</b>

Table 16: show interfaces ge Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> <li>• <b>Link partner speed</b>—Speed of the link partner.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports PAUSE on receive and transmit), <b>Asymmetric</b> (link partner supports PAUSE on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports PAUSE on both receive and transmit or PAUSE only on receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> <li>• <b>Destination slot</b>—FPC slot number.</li> </ul>	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels

Table 16: show interfaces ge Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>Protocol</b>	Protocol family.	<b>detail extensive none</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	<b>extensive</b>
<b>Local statistics</b>	Number and rate of bytes and packets destined to and from the switch.	<b>extensive</b>
<b>Transit statistics</b>	Number and rate of bytes and packets transiting the switch.	<b>extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .	<b>detail extensive none</b>
<b>Input Filters</b>	Names of any input filters applied to this interface.	<b>detail extensive</b>
<b>Output Filters</b>	Names of any output filters applied to this interface.	<b>detail extensive</b>
<b>Flags</b>	Information about protocol family flags.  If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.	<b>detail extensive</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>Flags</b>	Information about the address flags.	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interlace.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>



## Sample Output

### show interfaces

```

user@switch> show interfaces ge-0/0/9
Physical interface: ge-0/0/9, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 21
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
  Last flapped   : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 22)
  Flags: SNMP-Traps
  Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch
  Flags: None

```

### show interfaces brief

```

user@switch> show interfaces ge-0/0/9 brief
Physical interface: ge-0/0/9, Enabled, Physical link is Down
  Description: voice priority and tcp and icmp traffic rate-limiting filter at i
  ngress port
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None

Logical interface ge-0/0/9.0
  Flags: Device-Down SNMP-Traps Encapsulation: ENET2
  eth-switch

```

### show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
  Interface index: 193, SNMP ifIndex: 206, Generation: 196
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues

```

```

Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped    : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets

  0 best-effort      0              0              0
  1 assured-forw     0              0              0
  5 expedited-fo     0              0              0
  7 network-cont     0              0              0

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,

```

#### show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/9 detail
Physical interface: ge-0/0/9, Enabled, Physical link is Up
Interface index: 193, SNMP ifIndex: 206, Generation: 196
Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on ,
Auto-negotiation: Enabled,

```

```

Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped   : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0 0 bps
Output bytes  : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Egress queues: 8 supported, 4 in use
Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 assured-forw	0	0	0
5 expedited-fo	0	0	0
7 network-cont	0	0	0

```

Active alarms : None
Active defects : None

Logical interface ge-0/0/9.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
Flags: SNMP-Traps Encapsulation: ENET2
Bandwidth: 0
Traffic statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Local statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Transit statistics:
Input bytes   : 0 0 bps
Output bytes  : 0 0 bps
Input packets : 0 0 pps
Output packets: 0 0 pps
Protocol eth-switch, Generation: 146, Route table: 0
Flags: Is-Primary
Input Filters: f1,
Output Filters: f2,,,,

```

#### show interfaces extensive (Symmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
Interface index: 49164, SNMP ifIndex: 574, Generation: 142

```

```

Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
Last flapped   : 2011-02-25 00:45:03 UTC (22:42:48 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps
Input packets : 0          0 pps
Output packets: 0          0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort    0          0          0

    2 no-loss        0          0          0

    3 fcoe           0          0          0

    7 network-cont   0          0          0

Queue number:      Mapped forwarding classes
0                 best-effort
2                 no-loss
3                 fcoe
7                 network-control

Active alarms : LINK
Active defects : LINK
MAC statistics:
    Receive          Transmit
Total octets        0          0
Total packets       0          0
Unicast packets     0          0
Broadcast packets   0          0
Multicast packets   0          0
CRC/Align errors    0          0
FIFO errors         0          0
MAC control frames  0          0
MAC pause frames    0          0
Oversized frames    0
Jabber frames       0
Fragment frames     0

```

```

VLAN tagged frames          0
Code violations              0
MAC Priority Flow Control Statistics:
  Priority : 0                0          0
  Priority : 1                0          0
  Priority : 2                0          0
  Priority : 3                0          0
  Priority : 4                0          0
  Priority : 5                0          0
  Priority : 6                0          0
  Priority : 7                0          0
Filter statistics:
  Input packet count          0
  Input packet rejects        0
  Input DA rejects            0
  Input SA rejects            0
  Output packet count         0
  Output packet pad count     0
  Output packet error count   0
  CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth          Buffer Priority
Limit
    0 best-effort             75          750000000    75          0          low
none
    7 network-control         5           500000000    5           0          low
none
    8 mcast-be                15         1500000000   15          0          low
none
    11 mcast-nc               5           500000000    5           0          low
none

```

### show interfaces extensive (Asymmetric Flow Control and Autonegotiation Enabled)

```

user@switch> show interfaces ge-0/0/12 extensive
interface: ge-0/0/12, Enabled, Physical link is Down
  Interface index: 49164, SNMP ifIndex: 574, Generation: 142
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Duplex: Full-Duplex,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Configured-flow-control tx-buffers: off
rx-buffers: on
  Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:22:83:2a:d8:dc, Hardware address: 00:22:83:2a:d8:dc
  Last flapped   : 2011-02-25 00:45:03 UTC (22:42:48 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes : 0          0 bps
  Output bytes: 0          0 bps
  Input packets: 0        0 pps

```

```

Output packets:                                0                      0 pps
IPv6 transit statistics:
  Input bytes :                                0
  Output bytes :                               0
  Input packets:                               0
  Output packets:                              0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort              0                      0                      0
  2 no-loss                  0                      0                      0
  3 fcoe                     0                      0                      0
  7 network-cont             0                      0                      0

Queue number:      Mapped forwarding classes
  0                best-effort
  2                no-loss
  3                fcoe
  7                network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:
  Total octets              Receive      Transmit
  Total packets             0          0
  Unicast packets           0          0
  Broadcast packets         0          0
  Multicast packets         0          0
  CRC/Align errors          0          0
  FIFO errors               0          0
  MAC control frames        0          0
  MAC pause frames          0          0
  Oversized frames          0
  Jabber frames             0
  Fragment frames           0
  VLAN tagged frames        0
  Code violations           0
MAC Priority Flow Control Statistics:
  Priority : 0              0          0
  Priority : 1              0          0
  Priority : 2              0          0
  Priority : 3              0          0
  Priority : 4              0          0
  Priority : 5              0          0
  Priority : 6              0          0
  Priority : 7              0          0
Filter statistics:
  Input packet count        0
  Input packet rejects      0
  Input DA rejects          0
  Input SA rejects          0

```

```

Output packet count                                0
Output packet pad count                            0
Output packet error count                          0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link Partner:
Link mode: Full-duplex, Flow control: None, Remote fault: OK,
Link partner Speed: 1000 Mbps
Local resolution:
Flow control: enable PAUSE transmit and Disable PAUSE receive, Remote
fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
0 best-effort           75      750000000  75      0      low
none
7 network-control       5       500000000   5      0      low
none
8 mcast-be             15      1500000000  15      0      low
none
11 mcast-nc             5       500000000   5      0      low
none

```

#### show interfaces terse

```

user@switch> show interfaces ge-0/0/12 terse
Interface      Admin Link Proto  Local      Remote
ge-0/0/12      up    up

```

#### show interfaces terse (QFabric Systems)

```

user@switch> show interfaces node1:ge-0/0/0 terse
Physical interface: node1:ge-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled
Interface flags: Internal: 0x4000
CoS queues      : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped    : Never
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)

```

## show interfaces (GRE)


<b>Syntax</b>	<pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Display status information about the specified generic routing encapsulation (GRE) interface.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers and EX Series switches, the interface type is <b><i>gr-fpc/pic/port</i></b>.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified output level of interface information.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
	<p> <b>NOTE:</b> You can configure generic routing encapsulation (GRE) interfaces (gre-x/y/z) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the <i>Junos OS MPLS Applications Library for Routing Devices</i> and the <i>Junos OS, Release 14.1</i>.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (GRE) on page 156</a></p> <p><a href="#">show interfaces brief (GRE) on page 156</a></p> <p><a href="#">show interfaces detail (GRE) on page 156</a></p> <p><a href="#">show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 157</a></p> <p><a href="#">show interfaces extensive (GRE) on page 158</a></p>
<b>Output Fields</b>	Table 17 on page 153 lists the output fields for the <b>show interfaces (GRE)</b> command. Output fields are listed in the approximate order in which they appear.



Table 17: GRE show interfaces Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Type of interface.	All levels
<b>Link-level type</b>	Encapsulation used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Device Flags</b>	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface Flags</b>	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Input rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	<p>The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>

Table 17: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support.	detail extensive
Flags	<p>Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under <i>Common Output Fields Description</i>. describe general information about the logical interface.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>Reassemble-Pkts</b>—If the <b>Flags</b> field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation.</li> </ul>	All levels
IP-Header	<p>IP header of the logical interface. If the <b>tunnel key</b> statement is configured, this information is included in the <b>IP Header</b> entry.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>df</b>—If the <b>IP-Header</b> field includes this string immediately following the 16 bits of identification information (that is, if <b>:df:</b> displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation.</li> </ul>	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Copy-tos-to-outer-ip-header	<p>Status of type of service (ToS) bits in the GRE packet header:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel.</li> <li>• <b>Off</b>—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header.</li> </ul> <p><b>NOTE:</b> EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always <b>Off</b> in switch output.</p>	detail extensive
Gre keepalives configured	<p>Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel.</p> <p><b>NOTE:</b> EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always <b>Off</b> in switch output.</p>	detail extensive
Gre keepalives adjacency state	Status of the other end of the GRE tunnel: <b>Up</b> or <b>Down</b> . If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up.	detail extensive
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified

Table 17: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Rate of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> <li>• <b>Input rate</b>—Rate of bits and packets received on the interface.</li> <li>• <b>Output rate</b>—Rate of bits and packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	<b>detail extensive</b>
<b>Transit statistics</b>	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	<b>detail extensive none</b>
<b>Protocol</b>	Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , or <b>mpls</b> .	<b>detail extensive none</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>MTU</b>	MTU size on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## Sample Output

### show interfaces (GRE)

```

user@host> show interfaces gr-1/2/0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
  Flags: None
  Addresses, Flags: Is-Primary
    Local: 1.10.1.1

```

### show interfaces brief (GRE)

```

user@host> show interfaces gr-1/2/0 brief
Physical interface: gr-1/2/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface gr-1/2/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
  IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  inet 10.100.0.1/30
  mpls

```

### show interfaces detail (GRE)

```

user@host> show interfaces gr-1/2/0 detail
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26, Generation: 13
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags      : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0

```

```

Output packets:                0
Local statistics:
Input bytes :                  0
Output bytes :                  0
Input packets:                 0
Output packets:                0
Transit statistics:
Input bytes :                  0          0 bps
Output bytes :                  0          0 bps
Input packets:                 0          0 pps
Output packets:                0          0 pps
Protocol inet, MTU: 1476, Generation: 12, Route table: 0
Flags: None
Addresses, Flags: Is-Primary
Destination: Unspecified, Local: 1.10.1.1, Broadcast: Unspecified,
Generation: 15

```

### show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch

```

user@switch> show interfaces gr-2/0/15 detail
Physical interface: gr-2/0/15, Enabled, Physical link is Up
Interface index: 195, SNMP ifIndex: 846, Generation: 198
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:1f:12:38:0f:d2, Hardware address: 00:1f:12:38:0f:d2
Device flags    : Present Running
Interface flags: Point-To-Point SNMP-Traps
Statistics last cleared: 2011-09-14 17:43:15 UTC (00:00:18 ago)
Traffic statistics:
Input bytes :          5600636          0 bps
Output bytes :          5600636          0 bps
Input packets:          20007          0 pps
Output packets:          20007          0 pps
IPv6 transit statistics:
Input bytes :          0
Output bytes :          0
Input packets:          0
Output packets:          0

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 180.20.30.2:180.20.3:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
Input bytes :          5600886
Output bytes :          2881784
Input packets:          20010
Output packets:          10018
Local statistics:
Input bytes :          398
Output bytes :          264
Input packets:          5
Output packets:          3
Transit statistics:
Input bytes :          5600488          0 bps
Output bytes :          2881520          0 bps
Input packets:          20005          0 pps
Output packets:          10015          0 pps

```

```
Protocol inet, Generation: 159, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 90.90.90/24, Local: 90.90.90.10, Broadcast: 90.90.90.255,
  Generation: 144
```

```
Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)
```

```
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 160.20.40.2:160.20.30.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
```

```
Traffic statistics:
```

```
Input bytes :                260
Output bytes :              2880148
Input packets:                 4
Output packets:             10002
```

```
Local statistics:
```

```
Input bytes :                112
Output bytes :                 0
Input packets:                 2
Output packets:                 0
```

```
Transit statistics:
```

```
Input bytes :                148                0 bps
Output bytes :              2880148                0 bps
Input packets:                 2                0 pps
Output packets:             10002                0 pps
```

```
Protocol inet, Generation: 171, Route table: 0
```

```
Flags: None
```

```
Addresses, Flags: Is-Preferred Is-Primary
```

```
  Destination: 70.70.70/24, Local: 70.70.70.10, Broadcast: 70.70.70.255,
  Generation: 160
```

### [show interfaces extensive \(GRE\)](#)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 156](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 157](#).

## show interfaces queue

<b>Syntax</b>	<pre>show interfaces queue &lt;aggregate   remaining-traffic&gt; &lt;both-ingress-egress&gt; &lt;egress&gt; &lt;forwarding-class forwarding-class&gt; &lt;ingress&gt; &lt;interface-name interface-name&gt; &lt;l2-statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>both-ingress-egress</b>, <b>egress</b>, and <b>ingress</b> options introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p><b>l2-statistics</b> option introduced in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Display class-of-service (CoS) queue information for physical interfaces.
<b>Options</b>	<p><b>none</b>—Show detailed CoS queue statistics for all physical interfaces.</p> <p><b>aggregate</b>—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)</p> <p><b>both-ingress-egress</b>—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)</p> <p><b>egress</b>—(Optional) Display egress queue statistics.</p> <p><b>forwarding-class forwarding-class</b>—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.</p> <p><b>ingress</b>—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)</p> <p><b>interface-name interface-name</b>—(Optional) Show detailed CoS queue statistics for the specified interface.</p> <p><b>l2-statistics</b>—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles</p> <p><b>remaining-traffic</b>—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.</p>

### Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 18 on page 160](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 18 on page 160](#).

Table 18: Layer 2 Overhead and Transmitted Packets or Byte Counts

Protocol	Fragmentation		LFI
	First fragmentation	Second to <i>n</i> fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

#### Layer 2 Statistics—Fragmentation Overhead Calculation

##### MLPPP/MC-MLPPP Overhead details:

=====

##### Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

##### Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

##### MLFR (FRF15) Overhead details:

=====

##### Fragment 1:

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto               : 2 bytes
HDLC flag and FCS         : 4 bytes

```

##### Fragments 2 ...n :

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

##### MFR (FRF16) Overhead details:

=====



```

Fragment 1:
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  Inner proto          : 2 bytes
  HDLC flag and FCS    : 4 bytes

Fragments 2 ...n :
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  HDLC flag and FCS    : 4 bytes

```

## Overhead with LFI

```

MLPPP(Long & short sequence):
=====
  Outer PPP header : 4 bytes
  HDLC flag and FCS : 4 bytes

```

```

MLFR (FRF15):
=====
  Framereelay header : 2 bytes
  Control,NLPID      : 2 bytes
  HDLC flag and FCS  : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

**remaining-traffic**—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

## Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs), Modular Port Concentrators (MPCs), or Enhanced Queuing DPCs, rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



**NOTE:** For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

**Required Privilege Level**

view

**List of Sample Output**

[show interfaces queue \(Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC\) on page 167](#)  
[show interfaces queue \(Aggregated Ethernet on a T320 Router\) on page 168](#)  
[show interfaces queue \(Fast Ethernet on a J4300 Router\) on page 170](#)  
[show interfaces queue \(Gigabit Ethernet on a T640 Router\) on page 170](#)  
[show interfaces queue aggregate \(Gigabit Ethernet Enhanced DPC\) on page 171](#)  
[show interfaces queue \(Gigabit Ethernet IQ2 PIC\) on page 175](#)  
[show interfaces queue both-ingress-egress \(Gigabit Ethernet IQ2 PIC\) on page 178](#)  
[show interfaces queue ingress \(Gigabit Ethernet IQ2 PIC\) on page 180](#)  
[show interfaces queue egress \(Gigabit Ethernet IQ2 PIC\) on page 181](#)  
[show interfaces queue remaining-traffic \(Gigabit Ethernet Enhanced DPC\) on page 182](#)  
[show interfaces queue \(Channelized OC12 IQE Type 3 PIC in SONET Mode\) on page 185](#)  
[show interfaces queue \(QFX Series\) on page 195](#)  
[show interfaces queue l2-statistics \(lsq interface\) on page 196](#)  
[show interfaces queue lsq \(lsq-ifd\) on page 196](#)

**Output Fields** Table 19 on page 163 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

**Table 19: show interfaces queue Output Fields**

Field Name	Field Description
<b>Physical interface</b>	Name of the physical interface.
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.
<b>SNMP ifindex</b>	SNMP index number for the interface.
<b>Forwarding classes supported</b>	Total number of forwarding classes supported on the specified interface.
<b>Forwarding classes in use</b>	Total number of forwarding classes in use on the specified interface.
<b>Ingress queues supported</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
<b>Ingress queues in use</b>	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
<b>Output queues supported</b>	Total number of output queues supported on the specified interface.
<b>Output queues in use</b>	Total number of output queues in use on the specified interface.
<b>Egress queues supported</b>	Total number of egress queues supported on the specified interface.
<b>Egress queues in use</b>	Total number of egress queues in use on the specified interface.
<b>Queue counters (Ingress)</b>	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>
<b>Burst size</b>	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
The following output fields are applicable to both interface component and Packet Forwarding component in the <b>show interfaces queue</b> command:	
<b>Queue</b>	Queue number.
<b>Forwarding classes</b>	Forwarding class name.

Table 19: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
<b>Queued Packets</b>	<p>Number of packets queued to this queue.</p> <p><b>NOTE:</b> For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 161</a>.</p>
<b>Queued Bytes</b>	<p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 20 on page 166</a>.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see <a href="#">“Additional Information” on page 161</a>.</p>
<b>Transmitted Packets</b>	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the <b>Packet Forwarding Engine Chassis Queues</b> field) shows the prefragmentation values.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 159</a></p>
<b>Transmitted Bytes</b>	<p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see <a href="#">Table 20 on page 166</a>.</p> <p><b>NOTE:</b> On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p><b>NOTE:</b> For Layer 2 statistics, see <a href="#">“Overhead for Layer 2 Statistics” on page 159</a></p>
<b>Tail-dropped packets</b>	Number of packets dropped because of tail drop.
<b>RL-dropped packets</b>	<p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 161</a>.</p>
<b>RL-dropped bytes</b>	<p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see <a href="#">“Additional Information” on page 161</a>.</p>

Table 19: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP packets dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP packets dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP packets dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP packets dropped because of RED.</li> </ul> </li> <li>(J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority packets dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority packets dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority packets dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority packets dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see <a href="#">Table 20 on page 166</a>.</p> <ul style="list-style-type: none"> <li>(M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low, non-TCP</b>—Number of low-loss priority non-TCP bytes dropped because of RED.</li> <li><b>Low, TCP</b>—Number of low-loss priority TCP bytes dropped because of RED.</li> <li><b>High, non-TCP</b>—Number of high-loss priority non-TCP bytes dropped because of RED.</li> <li><b>High, TCP</b>—Number of high-loss priority TCP bytes dropped because of RED.</li> </ul> </li> <li>(J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> <li><b>Low</b>—Number of low-loss priority bytes dropped because of RED.</li> <li><b>Medium-low</b>—Number of medium-low loss priority bytes dropped because of RED.</li> <li><b>Medium-high</b>—Number of medium-high loss priority bytes dropped because of RED.</li> <li><b>High</b>—Number of high-loss priority bytes dropped because of RED.</li> </ul> </li> </ul> <p><b>NOTE:</b> Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Byte counts vary by interface hardware. [Table 20 on page 166](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 20 on page 166](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 20: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap.</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead.</li> </ul> <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN).</li> <li>RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted.</li> </ul>	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 20: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 478 bytes of Layer 3 packet.</li> </ul> <p>M Series routers:</p> <ul style="list-style-type: none"> <li>Queued: 478 bytes of Layer 3 packet.</li> <li>Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes</li> <li>RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet</li> </ul>	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	<p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	—
8-port 1G IQ2 and IQ2-E PICs		Transmitted: 478 bytes of Layer 3 packet.	

## Sample Output

### show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
  Transmitted:
    Packets          :          131300649          141751 pps
    Bytes            :          11287964840        99793248 bps
    Tail-dropped packets :              0              0 pps
    RL-dropped packets  :          205050862        602295 pps
    RL-dropped bytes    :          13595326612      327648832 bps
    RED-dropped packets :              0              0 pps
      Low              :              0              0 pps
      Medium-low       :              0              0 pps
      Medium-high      :              0              0 pps
      High             :              0              0 pps
    RED-dropped bytes   :              0              0 bps
      Low              :              0              0 bps
      Medium-low       :              0              0 bps
      Medium-high      :              0              0 bps
      High             :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets          :              0              0 pps
    Bytes            :              0              0 bps
```

### show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 33 Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :              5              0 pps
    Bytes            :             242              0 bps
  Transmitted:
    Packets          :              5              0 pps
    Bytes            :             242              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets :              0              0 pps
    RED-dropped bytes   :              0              0 bps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets          :          42603765          595484 pps
```



```

    Bytes                :          5453281920          609776496 bps
  Transmitted:
    Packets              :          42603765           595484 pps
    Bytes                :          5453281920          609776496 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets  :              0              0 pps
    RED-dropped bytes    :              0              0 bps
Queue: 2, Forwarding classes: ef1
  Queued:
    Packets              :              0              0 pps
    Bytes                :              0              0 bps
  Transmitted:
    Packets              :              0              0 pps
    Bytes                :              0              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets  :              0              0 pps
    RED-dropped bytes    :              0              0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets              :              45              0 pps
    Bytes                :             3930              0 bps
  Transmitted:
    Packets              :              45              0 pps
    Bytes                :             3930              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets  :              0              0 pps
    RED-dropped bytes    :              0              0 bps
Queue: 4, Forwarding classes: af11
  Queued:
    Packets              :              0              0 pps
    Bytes                :              0              0 bps
  Transmitted:
    Packets              :              0              0 pps
    Bytes                :              0              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets  :              0              0 pps
    RED-dropped bytes    :              0              0 bps
Queue: 5, Forwarding classes: ef11
  Queued:
    Packets              :              0              0 pps
    Bytes                :              0              0 bps
  Transmitted:
    Packets              :              0              0 pps
    Bytes                :              0              0 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets  :              0              0 pps
    RED-dropped bytes    :              0              0 bps
Queue: 6, Forwarding classes: af12
  Queued:
    Packets              :          31296413          437436 pps
    Bytes                :         4005940864         447935200 bps
  Transmitted:
    Packets              :          31296413          437436 pps
    Bytes                :         4005940864         447935200 bps
    Tail-dropped packets :              0              0 pps
    RED-dropped packets  :              0              0 pps
    RED-dropped bytes    :              0              0 bps
Queue: 7, Forwarding classes: nc2
  Queued:
    Packets              :              0              0 pps
    Bytes                :              0              0 bps

```

```

Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps

```

#### show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :                5240762                3404 pps
    Bytes            :            3020710354            15934544 bps
  Transmitted:
    Packets          :                5240762                3404 pps
    Bytes            :            3020710354            15934544 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
    Low               :                0                0 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                0                0 pps
    RED-dropped bytes  :                0                0 bps
    Low               :                0                0 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                0                0 pps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets          :                2480391                1650 pps
    Bytes            :            1304685666            6945704 bps
  Transmitted:
    Packets          :                2478740                1650 pps
    Bytes            :            1303817240            6945704 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                1651                0 pps
    Low               :                0                0 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                1651                0 pps
    RED-dropped bytes  :                868426                0 bps
    Low               :                0                0 pps
    Medium-low        :                0                0 pps
    Medium-high       :                0                0 pps
    High              :                868426                0 pps

```

#### show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:

```

```

Packets      :      13      0 pps
Bytes        :      622      0 bps
Transmitted:
Packets      :      13      0 pps
Bytes        :      622      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: af1
Queued:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Transmitted:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Transmitted:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Transmitted:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps

```

### show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets      :      148450735      947295 pps
Bytes        :      8016344944      409228848 bps
Transmitted:
Packets      :      76397439      487512 pps
Bytes        :      4125461868      210602376 bps
Tail-dropped packets : Not Available
RED-dropped packets :      72053285      459783 pps
Low          :      72053285      459783 pps
Medium-low   :      0      0 pps
Medium-high  :      0      0 pps
High         :      0      0 pps
RED-dropped bytes  :      3890877444      198626472 bps

```

```

Low : 3890877444 198626472 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 410278257 473940 pps
Bytes : 22156199518 204742296 bps
Transmitted:
Packets : 4850003 4033 pps
Bytes : 261900162 1742256 bps
Tail-dropped packets : Not Available
RED-dropped packets : 405425693 469907 pps
Low : 405425693 469907 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 21892988124 203000040 bps
Low : 21892988124 203000040 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```

```

Queued:
  Packets      :          76605230          485376 pps
  Bytes       :          5209211400        264044560 bps
Transmitted:
  Packets      :          76444631          484336 pps
  Bytes       :          5198235612        263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets :          160475          1040 pps
  Low         :          160475          1040 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :          10912300        565760 bps
  Low         :          10912300        565760 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Transmitted:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Tail-dropped packets : Not Available
RED-dropped packets :              0              0 pps
  Low         :              0              0 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :              0              0 bps
  Low         :              0              0 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :          4836136          3912 pps
  Bytes       :          333402032        2139056 bps
Transmitted:
  Packets      :          3600866          1459 pps
  Bytes       :          244858888        793696 bps
Tail-dropped packets : Not Available
RED-dropped packets :          1225034          2450 pps
  Low         :          1225034          2450 pps
  Medium-low  :              0              0 pps
  Medium-high :              0              0 pps
  High        :              0              0 pps
RED-dropped bytes  :          83302312        1333072 bps
  Low         :          83302312        1333072 bps
  Medium-low  :              0              0 bps
  Medium-high :              0              0 bps
  High        :              0              0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Transmitted:
  Packets      :              0              0 pps
  Bytes       :              0              0 bps
Tail-dropped packets : Not Available

```

RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	77059796	486384 pps
Bytes	:	3544750624	178989576 bps

Transmitted:

Packets	:	77059797	486381 pps
Bytes	:	3544750670	178988248 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	4846580	3934 pps
Bytes	:	222942680	1447768 bps

Transmitted:

Packets	:	4846580	3934 pps
Bytes	:	222942680	1447768 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps

```

      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps

```

### show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
  Transmitted:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps

```

```

RED-dropped bytes      :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets               :                7055              1 pps
  Bytes                 :            451552              512 bps
Transmitted:
  Packets               :                7055              1 pps
  Bytes                 :            451552              512 bps
Tail-dropped packets : Not Available
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets               :                1031              0 pps
  Bytes                 :            143292              0 bps
Transmitted:
  Packets               :                1031              0 pps
  Bytes                 :            143292              0 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Transmitted:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Transmitted:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets               :                77009             11 pps
  Bytes                 :            6894286             7888 bps
Transmitted:
  Packets               :                77009             11 pps
  Bytes                 :            6894286             7888 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps

```



```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          1031          0 pps
    Bytes        :        147328          0 bps
  Transmitted:
    Packets      :          1031          0 pps
    Bytes        :        147328          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    Low, non-TCP   :          0          0 pps
    Low, TCP       :          0          0 pps
    High, non-TCP  :          0          0 pps
    High, TCP      :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low, non-TCP   :          0          0 bps
    Low, TCP       :          0          0 bps
    High, non-TCP  :          0          0 bps
    High, TCP      :          0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    Low, non-TCP   :          0          0 pps
    Low, TCP       :          0          0 pps
    High, non-TCP  :          0          0 pps
    High, TCP      :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low, non-TCP   :          0          0 bps
    Low, TCP       :          0          0 bps
    High, non-TCP  :          0          0 bps
    High, TCP      :          0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    Low, non-TCP   :          0          0 pps
    Low, TCP       :          0          0 pps
    High, non-TCP  :          0          0 pps
    High, TCP      :          0          0 pps
    RED-dropped bytes :          0          0 bps
    Low, non-TCP   :          0          0 bps
    Low, TCP       :          0          0 bps
    High, non-TCP  :          0          0 bps
    High, TCP      :          0          0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :        94386         12 pps
    Bytes        :       13756799       9568 bps
  Transmitted:

```

Packets	:	94386	12 pps
Bytes	:	13756799	9568 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

### show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
  Forwarding classes: 8 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
  Queue: 0, Forwarding classes: best-effort
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                254                0 pps
      Bytes        :            16274                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 1, Forwarding classes: expedited-forwarding
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 2, Forwarding classes: assured-forwarding
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 3, Forwarding classes: network-control
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps

```

```

    RED-dropped bytes      :                0                0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                3                0 pps
    Bytes                  :               126                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                :             80564692            0 pps
    Bytes                  :          3383717100            0 bps
  Transmitted:
    Packets                :             80564692            0 pps
    Bytes                  :          3383717100            0 bps
    Tail-dropped packets   :                0                0 pps
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                :             80564685            0 pps
    Bytes                  :          3383716770            0 bps
  Transmitted:
    Packets                :             80564685            0 pps

```

```

Bytes : 3383716770 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Transmitted:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

#### show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 288 0 pps
Bytes : 18450 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

### show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
  Transmitted:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
  Transmitted:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
  Transmitted:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

#### show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
  Transmitted:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :      0      0 pps
    Low          :      0      0 pps

```

```

Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps

```

```

      High                :                0                0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                :            109355853            471736 pps
    Bytes                  :            7436199152          256627968 bps
  Transmitted:
    Packets                :            109355852            471736 pps
    Bytes                  :            7436198640          256627968 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low                    :                0                0 pps
    Medium-low             :                0                0 pps
    Medium-high            :                0                0 pps
    High                   :                0                0 pps
  RED-dropped bytes :
    Low                    :                0                0 bps
    Medium-low             :                0                0 bps
    Medium-high            :                0                0 bps
    High                   :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low                    :                0                0 pps
    Medium-low             :                0                0 pps
    Medium-high            :                0                0 pps
    High                   :                0                0 pps
  RED-dropped bytes :
    Low                    :                0                0 bps
    Medium-low             :                0                0 bps
    Medium-high            :                0                0 bps
    High                   :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low                    :                0                0 pps
    Medium-low             :                0                0 pps
    Medium-high            :                0                0 pps
    High                   :                0                0 pps
  RED-dropped bytes :
    Low                    :                0                0 bps
    Medium-low             :                0                0 bps
    Medium-high            :                0                0 bps
    High                   :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps

```



```

Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low            :                0                0 pps
  Medium-low     :                0                0 pps
  Medium-high    :                0                0 pps
  High           :                0                0 pps
RED-dropped bytes :                0                0 bps
  Low            :                0                0 bps
  Medium-low     :                0                0 bps
  Medium-high    :                0                0 bps
  High           :                0                0 bps

```

#### show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

  Interface index: 192, SNMP ifIndex: 1948

  Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

  Forwarding classes: 16 supported, 9 in use

  Egress queues: 8 supported, 8 in use

  Queue: 0, Forwarding classes: DEFAULT

  Queued:

    Packets          :                214886                13449 pps
    Bytes            :                9884756                5164536 bps

  Transmitted:

    Packets          :                214886                13449 pps
    Bytes            :                9884756                5164536 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
      Low            :                0                0 pps
      Medium-low     :                0                0 pps
      Medium-high    :                0                0 pps
      High           :                0                0 pps
    RED-dropped bytes :                0                0 bps
      Low            :                0                0 bps
      Medium-low     :                0                0 bps

```

Medium-high : 0 0 bps

High : 0 0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets : 0 0 pps

Bytes : 0 0 bps

Transmitted:

Packets : 0 0 pps

Bytes : 0 0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low : 0 0 pps

Medium-low : 0 0 pps

Medium-high : 0 0 pps

High : 0 0 pps

RED-dropped bytes : 0 0 bps

Low : 0 0 bps

Medium-low : 0 0 bps

Medium-high : 0 0 bps

High : 0 0 bps

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets : 0 0 pps

Bytes : 0 0 bps

Transmitted:

Packets : 0 0 pps

Bytes : 0 0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low : 0 0 pps

Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

```

        High                :                0                0 bps
Queue: 6, Forwarding classes: CLASS_V_OUTPUT
Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets  :                0                0 pps
    Low                   :                0                0 pps
    Medium-low            :                0                0 pps
    Medium-high           :                0                0 pps
    High                  :                0                0 pps
    RED-dropped bytes    :                0                0 bps
    Low                   :                0                0 bps
    Medium-low            :                0                0 bps
    Medium-high           :                0                0 bps
    High                  :                0                0 bps
Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS
Queued:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets  :                0                0 pps
    Low                   :                0                0 pps
    Medium-low            :                0                0 pps
    Medium-high           :                0                0 pps

```

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

Transmitted:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
---------	---	---	-------

```

Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
  Low          :          0          0 pps
  Medium-low   :          0          0 pps
  Medium-high  :          0          0 pps
  High         :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low          :          0          0 bps
  Medium-low   :          0          0 bps
  Medium-high  :          0          0 bps
  High         :          0          0 bps
Queue: 2, Forwarding classes: PRIVATE
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
  Low          :          0          0 pps
  Medium-low   :          0          0 pps
  Medium-high  :          0          0 pps
  High         :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low          :          0          0 bps
  Medium-low   :          0          0 bps

```

Medium-high : 0 0 bps

High : 0 0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets : 32843 0 pps

Bytes : 2641754 56 bps

Transmitted:

Packets : 32843 0 pps

Bytes : 2641754 56 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low : 0 0 pps

Medium-low : 0 0 pps

Medium-high : 0 0 pps

High : 0 0 pps

RED-dropped bytes : 0 0 bps

Low : 0 0 bps

Medium-low : 0 0 bps

Medium-high : 0 0 bps

High : 0 0 bps

Queue: 4, Forwarding classes: CLASS\_B\_OUTPUT

Queued:

Packets : 0 0 pps

Bytes : 0 0 bps

Transmitted:

Packets : 0 0 pps

Bytes : 0 0 bps

Tail-dropped packets : 0 0 pps

RED-dropped packets : 0 0 pps

Low : 0 0 pps



Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS\_C\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS\_V\_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS\_S\_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High : 0 0 bps

### show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes        : 0 0 bps
    Tail-dropped packets : Not Available

```

Total-dropped packets:	0	0 pps
Total-dropped bytes :	0	0 bps

### show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :           1           0 pps
    Bytes        :        1001           0 bps
  Transmitted:
    Packets      :           5           0 pps
    Bytes        :        1062           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :           1           0 pps
    Bytes        :        1500           0 bps
  Transmitted:
    Packets      :           6           0 pps
    Bytes        :        1573           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :           1           0 pps
    Bytes        :         512           0 bps
  Transmitted:
    Packets      :           3           0 pps
    Bytes        :         549           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
  Transmitted:
    Packets      :           0           0 pps
    Bytes        :           0           0 bps
    Tail-dropped packets :           0           0 pps
    RED-dropped packets :           0           0 pps
    RED-dropped bytes  :           0           0 bps
=====

```

### show interfaces queue lsq (lsq-ifd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0

```

```

Queue: 0, Forwarding classes: be
  Queued:
    Packets      :          55576          1206 pps
    Bytes        :       29622008       5145472 bps
  Transmitted:
    Packets      :          55576          1206 pps
    Bytes        :       29622008       5145472 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets  :          0          0 pps
    RL-dropped bytes    :          0          0 bps
    RED-dropped packets :          0          0 pps
      Low              :          0          0 pps
      Medium-low       :          0          0 pps
      Medium-high      :          0          0 pps
      High              :          0          0 pps
    RED-dropped bytes   :          0          0 bps
      Low              :          0          0 bps
      Medium-low       :          0          0 bps
      Medium-high      :          0          0 bps
      High              :          0          0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets  :          0          0 pps
    RL-dropped bytes    :          0          0 bps
    RED-dropped packets :          0          0 pps
      Low              :          0          0 pps
      Medium-low       :          0          0 pps
      Medium-high      :          0          0 pps
      High              :          0          0 pps
    RED-dropped bytes   :          0          0 bps
      Low              :          0          0 bps
      Medium-low       :          0          0 bps
      Medium-high      :          0          0 bps
      High              :          0          0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RL-dropped packets  :          0          0 pps
    RL-dropped bytes    :          0          0 bps
    RED-dropped packets :          0          0 pps
      Low              :          0          0 pps
      Medium-low       :          0          0 pps
      Medium-high      :          0          0 pps
      High              :          0          0 pps
    RED-dropped bytes   :          0          0 bps
      Low              :          0          0 bps
      Medium-low       :          0          0 bps
      Medium-high      :          0          0 bps
      High              :          0          0 bps
Queue: 3, Forwarding classes: nc

```

Queued:			
Packets	:	22231	482 pps
Bytes	:	11849123	2057600 bps
Transmitted:			
Packets	:	22231	482 pps
Bytes	:	11849123	2057600 bps
Tail-dropped packets	:	0	0 pps
RL-dropped packets	:	0	0 pps
RL-dropped bytes	:	0	0 bps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

## show interfaces xe

<b>Syntax</b>	<pre>show interfaces <i>device-name:type-fpc/pic/port</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;routing-instance (all   <i>instance-name</i>)&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.
<b>Options</b>	<p><b><i>device-name:type-fpc/pic/port</i></b>—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>routing-instance (all   <i>instance-name</i>)</b>—(Optional) Display the name of an individual routing instance or display all routing instances.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Monitoring Interface Status and Traffic on page 17</a></li> <li>• <a href="#">Troubleshooting Network Interfaces on page 17</a></li> <li>• <a href="#">Troubleshooting an Aggregated Ethernet Interface on page 37</a></li> <li>• <a href="#">Junos OS Network Interfaces Library for Routing Devices</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show interfaces on page 207</a></p> <p><a href="#">show interfaces (Asymmetric Flow Control) on page 208</a></p> <p><a href="#">show interfaces brief on page 208</a></p> <p><a href="#">show interfaces detail on page 208</a></p> <p><a href="#">show interfaces detail (Asymmetric Flow Control) on page 210</a></p> <p><a href="#">show interfaces extensive on page 211</a></p>

[show interfaces extensive \(Asymmetric Flow Control\) on page 213](#)

[show interfaces terse on page 215](#)

[show interfaces \(QFabric System\) on page 215](#)

**Output Fields** Table 21 on page 200 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

**Table 21: show interfaces xe Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface.	All levels
<b>Interface index</b>	Index number of the physical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>MTU</b>	Maximum transmission unit size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Duplex</b>	Duplex mode of the interface, either <b>Full-Duplex</b> or <b>Half-Duplex</b> .	All levels
<b>Loopback</b>	Loopback status: <b>Enabled</b> or <b>Disabled</b> . If loopback is enabled, type of loopback: <b>Local</b> or <b>Remote</b> .	All levels
<b>Source filtering</b>	Source filtering status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>LAN-PHY mode</b>	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
<b>Unidirectional</b>	Unidirectional link mode status for 10-Gigabit Ethernet interface: <b>Enabled</b> or <b>Disabled</b> for parent interface; <b>Rx-only</b> or <b>Tx-only</b> for child interfaces.	All levels
<b>Flow control</b>	Flow control status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>NOTE:</b> This field is only displayed if asymmetric flow control is not configured.		



Table 21: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Configured-flow-control</b>	<p>Configured flow control for the interface transmit buffers (<b>tx-buffers</b>) and receive buffers (<b>rx-buffers</b>):</p> <ul style="list-style-type: none"> <li><b>tx-buffers</b>—<b>On</b> if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer. <b>Off</b> if the interface is not configured to respond to received PAUSE messages.</li> <li><b>rx-buffers</b>—<b>On</b> if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer. <b>Off</b> if the interface is not configured to generate and send PAUSE messages.</li> </ul> <p><b>NOTE:</b> This field is only displayed if asymmetric flow control is configured.</p>	All levels
<b>Auto-negotiation</b>	Autonegotiation status: <b>Enabled</b> or <b>Disabled</b> .	All levels
<b>Remote-fault</b>	<p>Remote fault status:</p> <ul style="list-style-type: none"> <li><b>Online</b>—Autonegotiation is manually configured as online.</li> <li><b>Offline</b>—Autonegotiation is manually configured as offline.</li> </ul>	All levels
<b>Device flags</b>	Information about the physical device.	All levels
<b>Interface flags</b>	Information about the interface.	All levels
<b>Link flags</b>	Information about the link.	All levels
<b>Wavelength</b>	Configured wavelength, in nanometers (nm).	All levels
<b>Frequency</b>	Frequency associated with the configured wavelength, in terahertz (THz).	All levels
<b>CoS queues</b>	Number of CoS queues configured.	<b>detail extensive none</b>
<b>Schedulers</b>	Number of CoS schedulers configured.	<b>extensive</b>
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive none</b>
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive none</b>
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago)</b> .	<b>detail extensive none</b>
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>

Table 21: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul> <p><b>NOTE:</b> The bandwidth bps counter is not enabled.</p>	<b>detail extensive</b>
<b>Input errors</b>	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle.</li> <li>• <b>L3 incompletes</b>—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the <b>ignore-l3-incompletes</b> statement.</li> <li>• <b>L2 channel errors</b>—Number of times the software did not find a valid logical interface for an incoming frame.</li> <li>• <b>L2 mismatch timeouts</b>—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>

Table 21: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Output errors</b>	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Collisions</b>—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug.</li> <li>• <b>Aged packets</b>—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware.</li> <li>• <b>FIFO errors</b>—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>HS link CRC errors</b>—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.</li> <li>• <b>MTU errors</b>—Number of packets whose size exceeded the MTU of the interface.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Egress queues</b>	Total number of egress queues supported on the specified interface.	<b>detail extensive</b>
<b>Queue counters (Egress)</b>	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>detail extensive</b>
<b>Queue Number</b>	The CoS queue number and the forwarding classes mapped to the queue number. The <b>Mapped forwarding class</b> column lists the forwarding classes mapped to each CoS queue.	<b>detail extensive</b>
<b>Ingress queues</b>	Total number of ingress queues supported on the specified interface.	<b>extensive</b>
<b>Queue counters (Ingress)</b>	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>extensive</b>

Table 21: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Active alarms and Active defects</b>	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value <b>None</b> or <b>Link</b>.</p> <ul style="list-style-type: none"> <li>• <b>None</b>—There are no active defects or alarms.</li> <li>• <b>Link</b>—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.</li> </ul>	<b>detail extensive none</b>
<b>PCS statistics</b>	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	<b>detail extensive</b>
<b>MAC statistics</b>	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> <li>• <b>Total octets and total packets</b>—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type.</li> <li>• <b>Unicast packets, Broadcast packets, and Multicast packets</b>—Number of unicast, broadcast, and multicast packets.</li> <li>• <b>CRC/Align errors</b>—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error).</li> <li>• <b>FIFO error</b>—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning.</li> <li>• <b>MAC control frames</b>—Number of MAC control frames.</li> <li>• <b>MAC pause frames</b>—Number of MAC control frames with <b>pause</b> operational code.</li> <li>• <b>Oversized frames</b>—Number of packets that exceeds the configured MTU.</li> <li>• <b>Jabber frames</b>—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms.</li> <li>• <b>Fragment frames</b>—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runs (which are normal occurrences caused by collisions) and noise hits are counted.</li> <li>• <b>VLAN tagged frames</b>—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0.</li> <li>• <b>Code violations</b>—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error."</li> </ul>	<b>extensive</b>
<b>Filter statistics</b>	Receive and Transmit statistics reported by the PIC's MAC address filter subsystem.	<b>extensive</b>

Table 21: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> <li>• <b>Negotiation status:</b> <ul style="list-style-type: none"> <li>• <b>Incomplete</b>—Ethernet interface has the speed or link mode configured.</li> <li>• <b>No autonegotiation</b>—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.</li> <li>• <b>Complete</b>—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> </ul> </li> <li>• <b>Link partner status</b>—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.</li> <li>• <b>Link partner:</b> <ul style="list-style-type: none"> <li>• <b>Link mode</b>—Depending on the capability of the attached Ethernet device, either <b>Full-duplex</b> or <b>Half-duplex</b>.</li> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is <b>None</b>. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive).</li> <li>• <b>Remote fault</b>—Remote fault information from the link partner—<b>Failure</b> indicates a receive link error. <b>OK</b> indicates that the link partner is receiving. <b>Negotiation error</b> indicates a negotiation error. <b>Offline</b> indicates that the link partner is going offline.</li> </ul> </li> <li>• <b>Local resolution:</b> <ul style="list-style-type: none"> <li>• <b>Flow control</b>—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are <b>Symmetric</b> (link partner supports <b>PAUSE</b> on receive and transmit), <b>Asymmetric</b> (link partner supports <b>PAUSE</b> on transmit), and <b>Symmetric/Asymmetric</b> (link partner supports both <b>PAUSE</b> on receive and transmit or only <b>PAUSE</b> receive). For asymmetric <b>PAUSE</b>, shows if the <b>PAUSE</b> transmit and <b>PAUSE</b> receive states on the interface are <b>enable</b> or <b>disable</b>.</li> <li>• <b>Remote fault</b>—Remote fault information. <b>Link OK</b> (no error detected on receive), <b>Offline</b> (local interface is offline), and <b>Link Failure</b> (link error detected on receive).</li> </ul> </li> </ul>	extensive

Table 21: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Packet Forwarding Engine configuration</b>	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> <li><b>Destination slot</b>—FPC slot number.</li> <li><b>CoS transmit queue</b>—Queue number and its associated user-configured forwarding class name.</li> <li><b>Bandwidth %</b>—Percentage of bandwidth allocated to the queue.</li> <li><b>Bandwidth bps</b>—Bandwidth allocated to the queue (in bps).</li> <li><b>Buffer %</b>—Percentage of buffer space allocated to the queue.</li> <li><b>Buffer usec</b>—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time.</li> <li><b>Priority</b>—Queue priority: <b>low</b> or <b>high</b>.</li> <li><b>Limit</b>—Displayed if rate limiting is configured for the queue. Possible values are <b>none</b> and <b>exact</b>. If <b>exact</b> is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If <b>none</b> is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.</li> </ul>	<b>extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Index number of the logical interface, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP interface index number for the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface.	All levels
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>Protocol</b>	Protocol family.	<b>detail extensive none</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	<b>detail extensive</b>
<b>IPv6 transit statistics</b>	If IPv6 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	<b>extensive</b>
<b>Local statistics</b>	Number and rate of bytes and packets destined to and from the switch.	<b>extensive</b>
<b>Transit statistics</b>	Number and rate of bytes and packets transiting the switch.	<b>extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Route table in which the logical interface address is located. For example, <b>0</b> refers to the routing table inet.0.	<b>detail extensive none</b>

Table 21: show interfaces xe Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Input Filters</b>	Names of any input filters applied to this interface.	<b>detail extensive</b>
<b>Output Filters</b>	Names of any output filters applied to this interface.	<b>detail extensive</b>
<b>Flags</b>	Information about protocol family flags.  If unicast Reverse Path Forwarding (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled.	<b>detail extensive</b>
<b>Addresses, Flags</b>	Information about the address flags.	<b>detail extensive none</b>
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>Flags</b>	Information about the address flag.	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## Sample Output

### show interfaces

```

user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped  : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0

```

```
Output packets: 0
Protocol eth-switch, MTU: 0
Flags: Trunk-Mode
```

### show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode
```

### show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

Logical interface xe-0/0/1.0
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch
```

### show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
```



```

Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Egress queues: 12 supported, 9 in use
Queue counters:
  Queued packets  Transmitted packets  Dropped packets

  0 best-effort 0 0 0
  1 fc7 0 0 0
  2 no-loss 0 0 0
  3 fcoe 0 0 0
  4 fc4 0 0 0
  5 fc5 0 0 0
  6 fc6 0 0 0
  7 network-cont 0 0 0
  8 mcast 0 0 0

Queue number: Mapped forwarding classes
0 best-effort
1 fc7
2 no-loss
3 fcoe
4 fc4
5 fc5
6 fc6
7 network-control
8 mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps

```

```

Input packets:          0          0 pps
Output packets:         0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces detail (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 detail
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:50 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
  Traffic statistics:
    Input bytes :          0          0 bps
    Output bytes :          0          0 bps
    Input packets:          0          0 pps
    Output packets:          0          0 pps
  IPv6 transit statistics:
    Input bytes :          0
    Output bytes :          0
    Input packets:          0
    Output packets:          0
  Egress queues: 12 supported, 9 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort          0              0              0
    1 fc7                 0              0              0
    2 no-loss              0              0              0
    3 fcoe                 0              0              0
    4 fc4                  0              0              0
    5 fc5                  0              0              0
    6 fc6                  0              0              0
    7 network-cont        0              0              0
    8 mcast                0              0              0

  Queue number:      Mapped forwarding classes
    0                best-effort
    1                fc7
    2                no-loss
    3                fcoe
    4                fc4
    5                fc5
    6                fc6

```

```

7          network-control
8          mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces extensive

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Flow control: Disabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

```

0 best-effort	0	0	0
1 fc7	0	0	0
2 no-loss	0	0	0
3 fcoe	0	0	0
4 fc4	0	0	0
5 fc5	0	0	0
6 fc6	0	0	0
7 network-cont	0	0	0
8 mcast	0	0	0

Queue number:            Mapped forwarding classes

0	best-effort
1	fc7
2	no-loss
3	fcoe
4	fc4
5	fc5
6	fc6
7	network-control
8	mcast

Active alarms : None

Active defects : None

MAC statistics:

	Receive	Transmit
Total octets	0	0
Total packets	0	0
Unicast packets	0	0
Broadcast packets	0	0
Multicast packets	0	0
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	

MAC Priority Flow Control Statistics:

Priority : 0	0	0
Priority : 1	0	0
Priority : 2	0	0
Priority : 3	0	0
Priority : 4	0	0
Priority : 5	0	0
Priority : 6	0	0
Priority : 7	0	0

Filter statistics:

Input packet count	0	
Input packet rejects	0	
Input DA rejects	0	
Input SA rejects	0	
Output packet count		0

```

Output packet pad count                                0
Output packet error count                              0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
      %      bps      %      usec
0 best-effort      75      7500000000      75      0      low
none
7 network-control      5      500000000      5      0      low
none
8 mcast      20      2000000000      20      0      low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0      0 bps
  Output bytes : 0      0 bps
  Input packets: 0      0 pps
  Output packets: 0      0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces extensive (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
  Input bytes : 0      0 bps
  Output bytes : 0      0 bps
  Input packets: 0      0 pps
  Output packets: 0      0 pps
IPv6 transit statistics:
  Input bytes : 0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          0                0                0
1 fc7                  0                0                0
2 no-loss              0                0                0
3 fcoe                 0                0                0
4 fc4                  0                0                0
5 fc5                  0                0                0
6 fc6                  0                0                0
7 network-cont         0                0                0
8 mcast                0                0                0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets    0            0
Unicast packets  0            0
Broadcast packets 0            0
Multicast packets 0            0
CRC/Align errors 0            0
FIFO errors       0            0
MAC control frames 0            0
MAC pause frames  0            0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
MAC Priority Flow Control Statistics:
Priority : 0       0            0
Priority : 1       0            0

```

```

Priority : 2          0          0
Priority : 3          0          0
Priority : 4          0          0
Priority : 5          0          0
Priority : 6          0          0
Priority : 7          0          0
Filter statistics:
Input packet count    0
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue    Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort         75    7500000000    75      0    low    none
7 network-control     5     500000000     5      0    low    none
8 mcast               20    2000000000    20      0    low    none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0          0 bps
Output bytes : 0          0 bps
Input packets: 0          0 pps
Output packets: 0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

### show interfaces terse

```

user@switch> show interfaces xe-0/0/1 terse
Interface      Admin Link Proto  Local      Remote

xe-0/0/1       up    up
xe-0/0/1.0     up    up    eth-switch

```

### show interfaces (QFabric System)

```

user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```

Interface flags: Internal: 0x4000  
CoS queues : 8 supported, 8 maximum usable queues  
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00  
Last flapped : Never  
Input rate : 0 bps (0 pps)  
Output rate : 0 bps (0 pps)



## CHAPTER 10

# LAGs and LACP Operational Commands

- `show lacp interfaces`
- `show lacp statistics interfaces` (View)

## show lacp interfaces

---

<b>Syntax</b>	<code>show lacp interfaces</code> <code>&lt;interface-name&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.
<b>Options</b>	<code>none</code> —Display LACP information for all interfaces.  <code>interface-name</code> —(Optional) Display LACP information for the specified interface: <ul style="list-style-type: none"><li>• Aggregated Ethernet—<code>aex</code></li><li>• Gigabit Ethernet—<code>ge-fpc/pic/port</code></li><li>• 10-Gigabit Ethernet—<code>xe-fpc/pic/port</code></li></ul>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li><li>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i></li><li>• <i>Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch</i></li><li>• <i>Configuring Aggregated Ethernet Links (CLI Procedure)</i></li><li>• <a href="#">Configuring Link Aggregation on page 33</a></li><li>• <i>Configuring Aggregated Ethernet LACP (CLI Procedure)</i></li><li>• <a href="#">Configuring Aggregated Ethernet LACP on page 32</a></li><li>• <i>Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)</i></li><li>• <i>Understanding Aggregated Ethernet Interfaces and LACP</i></li><li>• <a href="#">Understanding Aggregated Ethernet Interfaces and LACP on page 29</a></li><li>• <a href="#">Junos OS Interfaces Fundamentals Configuration Guide</a></li></ul>
<b>List of Sample Output</b>	<a href="#">show lacp interfaces (EX Series Switches) on page 220</a> <a href="#">show lacp interfaces (QFX Series) on page 221</a>
<b>Output Fields</b>	<a href="#">Table 22 on page 219</a> lists the output fields for the <code>show lacp interfaces</code> command. Output fields are listed in the approximate order in which they appear.

Table 22: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated Ethernet interface name.
LACP State	<p>LACP state information for each aggregated Ethernet interface:</p> <ul style="list-style-type: none"> <li>For a child interface configured with the <b>force-up</b> statement, LACP state displays <b>FUP</b> along with the interface name.</li> <li><b>Role</b>—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> <li><b>Actor</b>—Local device participating in the LACP negotiation.</li> <li><b>Partner</b>—Remote device participating in the LACP negotiation.</li> </ul> </li> <li><b>Exp</b>—Expired state. <b>Yes</b> indicates that the actor or partner is in an expired state. <b>No</b> indicates that the actor or partner is not in an expired state.</li> <li><b>Def</b>—Default. <b>Yes</b> indicates that the actor's receive machine is using the default operational partner information, which is administratively configured for the partner. <b>No</b> indicates that the operational partner information in use has been received in an LACP PDU.</li> <li><b>Dist</b>—Distribution of outgoing frames. <b>No</b> indicates that the distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is <b>Yes</b>.</li> <li><b>Col</b>—Collection of incoming frames. <b>Yes</b> indicates that the collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is <b>No</b>.</li> <li><b>Syn</b>—Synchronization. If the value is <b>Yes</b>, the link is considered to be synchronized. The link has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is <b>No</b>, the link is not synchronized. The link is currently not in the right aggregation.</li> <li><b>Aggr</b>—Ability of the aggregation port to aggregate (<b>Yes</b>) or to operate only as an individual link (<b>No</b>).</li> <li><b>Timeout</b>—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or a fast transmission rate, depending upon the expressed LACP timeout preference (<b>Long Timeout</b> or <b>Short Timeout</b>).</li> <li><b>Activity</b>—Actor's or partner's port activity. <b>Passive</b> indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is <b>Active</b>. <b>Active</b> indicates the port's preference to participate in the protocol regardless of the partner's control value.</li> </ul>

Table 22: show lacp interfaces Output Fields (*continued*)

Field Name	Field Description
LACP Protocol	<p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> <li>Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.</li> <li><b>Receive State</b>—One of the following values: <ul style="list-style-type: none"> <li><b>Current</b>—The state machine receives an LACP PDU and enters the <b>Current</b> state.</li> <li><b>Defaulted</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires a second time, the state machine enters the <b>Defaulted</b> state.</li> <li><b>Expired</b>—If no LACP PDU is received before the timer for the <b>Current</b> state expires once, the state machine enters the <b>Expired</b> state.</li> <li><b>Initialize</b>—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the <b>Initialize</b> state.</li> <li><b>LACP Disabled</b>—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to <b>LACP Disabled</b>. This state is similar to the <b>Defaulted</b> state, except that the port is forced to operate as an individual port.</li> <li><b>Port Disabled</b>—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the <b>Port Disabled</b> state.</li> </ul> </li> <li><b>Transmit State</b>—Transmit state of the state machine. The transmit state is one of the following values: <ul style="list-style-type: none"> <li><b>Fast periodic</b>—Periodic transmissions are enabled at a fast transmission rate.</li> <li><b>No periodic</b>—Periodic transmissions are disabled.</li> <li><b>Periodic timer</b>—Transitory state entered when the periodic timer expires.</li> <li><b>Slow periodic</b>—Periodic transmissions are enabled at a slow transmission rate.</li> </ul> </li> <li><b>Mux State</b>—State of the multiplexer state machine for the aggregation port. The state is one of the following values: <ul style="list-style-type: none"> <li><b>Attached</b>—The multiplexer state machine initiates the process of attaching the port to the selected aggregator.</li> <li><b>Collecting—Yes</b> indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. <b>No</b> indicates the receive function of this link is not enabled.</li> <li><b>Collecting distributing</b>—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.</li> <li><b>Detached</b>—Process of detaching the port from the aggregator is in progress.</li> <li><b>Distributing—Yes</b> indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames can be passed down from the aggregator's distribution function for transmission. <b>No</b> indicates the transmit function of this link is not enabled.</li> <li><b>Waiting</b>—The multiplexer state machine is in a holding process, awaiting an outcome.</li> </ul> </li> </ul>

## Sample Output

### show lacp interfaces (EX Series Switches)

```

user@switch> show lacp interfaces ae5
Aggregated interface: ae5
  LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
    xe-2/0/7      Actor  No   No   Yes   Yes  Yes   Yes    Fast    Active
    xe-2/0/7      Partner No   No   Yes   Yes  Yes   Yes    Fast    Passive

```

xe-4/0/7	Actor	No	No	No	No	No	Yes	Fast	Active
xe-4/0/7	Partner	No	No	No	Yes	Yes	Yes	Fast	Passive

LACP protocol:	Receive State	Transmit State	Mux State
xe-2/0/7(Active)	Current	Fast periodic	Collecting distributing
xe-34/0/7(Standby)	Current	Fast periodic	Waiting

### show lacp interfaces (QFX Series)

```
user@switch> show lacp interfaces nodegroup1:ae0 extensive
```

```
Aggregated interface: nodegroup1:ae0
```

LACP state:	Role	Exp	Def	Dist	Col	Syn	Aggr	Timeout	Activity
node1:xe-0/0/1FUP	Actor		No	Yes	No	No	No	Yes	Fast
Active									
node1xe-0/0/1FUP	Partner		No	Yes	No	No	No	Yes	Fast
Passive									
node2:xe-0/0/2	Actor		No	Yes	No	No	No	Yes	Fast
Active									
node2:xe-0/0/2	Partner		No	Yes	No	No	No	Yes	Fast
Passive									

LACP protocol:	Receive State	Transmit State	Mux State
node1:xe-0/0/1FUP	Current	Fast periodic	Collecting
distributing			
node2:xe-0/0/2	Current	Fast periodic	Collecting
distributing			
node1:xe-0/0/1 (active)	Current	Fast periodic	Collecting
distributing			
node2:xe-0/0/2 (standby)	Current	Fast periodic	WAITING

## show lacp statistics interfaces (View)

<b>Syntax</b>	<b>show lacp statistics interfaces</b> <i>interface-name</i>
<b>Release Information</b>	Command modified in Release 10.2 of Junos OS. Command introduced in Release 11.1 of Junos OS for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed.
<b>Options</b>	<i>interface-name</i> —(Optional) Name of an interface.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Link Aggregation on page 33</a></li> <li>• <a href="#">Verifying the Status of a LAG Interface on page 36</a></li> <li>• <a href="#">Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 36</a></li> <li>• <a href="#">Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch</a></li> <li>• <a href="#">Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show lacp statistics interfaces on page 224</a> <a href="#">show lacp statistics interfaces (QFX Series) on page 224</a> <a href="#">show lacp statistics interfaces (QFabric Systems) on page 224</a>
<b>Output Fields</b>	Table 23 on page 223 lists the output fields for the <b>show lacp statistics interfaces</b> command. Output fields are listed in the approximate order in which they appear.

**Table 23: show lacp statistics interfaces Output Fields**

Field Name	Field Description
Aggregated interface	Aggregated interface value.
LACP Statistics	<p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> <li>• <b>LACP Rx</b>—LACP received counter that increments for each normal hello.</li> <li>• <b>LACP Tx</b>—LACP transmit counter that increments for each normal hello.</li> <li>• <b>Unknown Rx</b>—Number of unrecognized packet errors logged.</li> <li>• <b>Illegal Rx</b>—Number of invalid packets received.</li> </ul>

## Sample Output

### show lacp statistics interfaces

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
ge-2/0/0              1352        2035          0                0
ge-2/0/1              1352        2056          0                0
ge-2/2/0              1352        2045          0                0
ge-2/2/1              1352        2043          0                0
```

### show lacp statistics interfaces (QFX Series)

```
user@host> show lacp statistics interfaces ae0
Aggregated interface: ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
xe-0/0/2              1352        2035          0                0
xe-0/0/3              1352        2056          0                0
```

### show lacp statistics interfaces (QFabric Systems)

```
user@host> show lacp statistics interfaces nodegroup1:ae0
Aggregated interface: nodegroup1:ae0
LACP Statistics:      LACP Rx      LACP Tx      Unknown Rx      Illegal Rx
node1:xe-0/0/2        1352        2035          0                0
node2:xe-0/0/3        1352        2056          0                0
```