Junos® OS for EX Series Ethernet Switches

Interfaces for EX Series Switches

Release

14.1X53

Modified: 2015-05-27

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

Junos® OS for EX Series Ethernet Switches Interfaces for EX Series Switches
Release 14.1X53

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

Table of Contents

Part 2 Configuration

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

Table 1 on page xix 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 xix defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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

Convention	Description	Examples
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention 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.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.

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

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

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

CHAPTER 1

Interfaces Overview

EX Series Switches Interfaces Overview

Juniper Networks EX Series Ethernet Switches 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 Interfaces Fundamentals Configuration Guide*.

For information about interface-naming conventions on EX Series switches, see "Understanding Interface Naming Conventions on EX Series Switches" on page 6.

This topic describes:

Network Interfaces

Network interfaces connect to the network and carry network traffic. Table 3 on page 4 lists the types of network interfaces supported on EX Series switches.

Table 3: Network Interface Types and Purposes

Type	Purpose
Aggregated Ethernet interfaces	All EX Series switches allow you to group Ethernet interfaces at the physical layer to form a single link layer interface, also known as a *link aggregation group (LAG)* or *bundle*. These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.
LAN access interfaces	Use these EX Series switch interfaces to connect a personal computer, laptop, file server, or printer to the network. When you power on an EX Series switch and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports. The default configuration also enables autonegotiation for both speed and link mode.
Power over Ethernet (PoE) interfaces	EX Series switches provide PoE network ports with various switch models. These ports can be used to connect voice over IP (VoIP) telephones, wireless access points, video cameras, and point-of-sale devices to safely receive power from the same access ports that are used to connect personal computers to the network. PoE interfaces are enabled by default in the factory configuration.
Trunk interfaces	EX Series access switches can be connected to a distribution switch or customer-edge (CE) switches or routers. To use a port for this type of connection, you must explicitly configure the network interface for trunk mode. The interfaces from the distribution switch or CE switch to the access switches must also be configured for trunk mode.

Special Interfaces

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

Table 4: Special Interface Types and Purposes

Type	Purpose
Console port	Each EX Series switch has a serial port, labeled **CON** or **CONSOLE**, for connecting tty-type terminals to the switch using standard PC-type tty cables. 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. On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, or an EX4500 Virtual Chassis, you can access the master and configure all members of the Virtual Chassis through any member's console port. For more information about the console port in a Virtual Chassis, see *Understanding Global Management of a Virtual Chassis*.
Loopback	All EX Series switches have this 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	The Juniper Networks Junos operating system (Junos OS) for EX Series switches automatically creates the switch's management Ethernet interface, **me0**. The management Ethernet interface provides an out-of-band method for connecting to the switch. To use **me0** as a management port, you must configure its logical port, **me0.0**, with a valid IP address. You can connect to the management interface over the network using utilities such as SSH or Telnet. SNMP can use the management interface to gather statistics from the switch. (The management interface **me0** is analogous to the **fxp0** interfaces on routers running Junos OS.)

Table 4: Special Interface Types and Purposes *(continued)*

<table>
<tr><th>Type</th><th>Purpose</th></tr>
<tr><td>Integrated Routing and Bridging (IRB) Interface or Routed VLAN Interface (RVI)</td><td>EX Series switches use an integrated routing and bridging (IRB) interface or Routed VLAN Interface (RVI) to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.

The IRB interface or RVI functions as a logical router, eliminating the need for having both a switch and a router. These interfaces must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed from.</td></tr>
<tr><td>Virtual Chassis port (VCP) interfaces</td><td>Virtual Chassis ports (VCPs) are used to interconnect switches in a Virtual Chassis:
EX3300 switches—Port 2 and port 3 of the SFP+ uplink ports are preconfigured as VCPs and can be used to interconnect up to six EX3300 switches in an EX3300 Virtual Chassis. See Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure).
EX4200 and EX4500 switches—Each EX4200 switch or each EX4500 switch with a Virtual Chassis module installed has two dedicated VCPs on its rear panel. These ports can be used to interconnect up to ten EX4200 switches in an EX4200 Virtual Chassis, up to ten EX4500 switches in an EX4500 Virtual Chassis, and up to ten switches in a mixed EX4200 and EX4500 Virtual Chassis. When you power on switches that are interconnected in this manner, the software automatically configures the VCP interfaces for the dedicated ports that have been interconnected. These VCP interfaces are not configurable or modifiable. See Understanding the High-Speed Interconnection of the Dedicated Virtual Chassis Ports Connecting EX4200, EX4500, and EX4550 Member Switches.
You can also interconnect EX4200 and EX4500 switches by using uplink module ports. Using uplink ports allows you to connect switches over longer distances than you can by using the dedicated VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure) or Setting an Uplink Port as a Virtual Chassis Port on an EX4500 or EX4550 Switch (CLI Procedure).
EX4300 switches—All QSFP+ ports are configured as VCPs, by default. See Understanding EX4300 Virtual Chassis
You can also interconnect EX4300 switches into a Virtual Chassis by using SFP+ uplink module ports as VCPs. Using uplink ports as VCPs allows you to connect switches over longer distances than you can by using the QSFP+ ports as VCPs. To use the uplink ports as VCPs, you must explicitly configure the uplink module ports on the members you want to connect as VCPs. See Setting an Uplink Port on an EX Series Switch as a Virtual Chassis Port (CLI Procedure).
EX8200 switches—EX8200 switches can be connected to an XRE200 External Routing Engine to create an EX8200 Virtual Chassis. The XRE200 External Routing Engine has dedicated VCPs that connect to ports on the internal Routing Engines of the EX8200 switches and can connect to another XRE200 External Routing Engine for redundancy. These ports require no configuration.
You can also connect two members of an EX8200 Virtual Chassis so that they can exchange Virtual Chassis Control Protocol (VCCP) traffic. To do so, you explicitly configure network ports on the EX8200 switches as VCPs. See Understanding Virtual Chassis Ports in an EX8200 Virtual Chassis.</td></tr>
<tr><td>Virtual management Ethernet (VME) interface</td><td>EX3300, EX4200, EX4300, and EX4500 switches have a VME interface. This is a logical interface that is used for Virtual Chassis configurations and allows you to manage all the members of the Virtual Chassis through the master. For more information about the VME interface, see Understanding Global Management of a Virtual Chassis.

EX8200 switches do not use a VME interface. An EX8200 Virtual Chassis is managed through the management Ethernet (me0) interface on the XRE200 External Routing Engine.</td></tr>
</table>

Related Documentation

- *EX2200 Switches Hardware Overview*
- *EX3200 Switches Hardware Overview*
- *EX3300 Switches Hardware Overview*
- *EX4200 Switches Hardware Overview*
- *EX4300 Switches Hardware Overview*
- *EX4500 Switches Hardware Overview*
- *EX6210 Switch Hardware Overview*
- *EX8208 Switch Hardware Overview*
- *EX8216 Switch Hardware Overview*
- *XRE200 External Routing Engine Hardware Overview*
- *Understanding PoE on EX Series Switches*
- Understanding Aggregated Ethernet Interfaces and LACP on page 8
- Understanding Layer 3 Subinterfaces on page 11

Understanding Interface Naming Conventions on EX Series Switches

Juniper Networks EX Series Ethernet Switches use a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos operating system (Junos OS). This topic provides brief information about the naming conventions used for interfaces on EX Series switches. For additional information, see the *Junos OS Network Interfaces Configuration Guide*.

This topic describes:

Physical Part of an Interface Name

Network interfaces in Junos OS are specified as follows:

type-fpc / pic / port

EX Series switches apply this convention as follows:

- *type*—EX Series interfaces use the following media types:
  - **ge**—Gigabit Ethernet interface
  - **xe**—10 Gigabit Ethernet interface

    - **et**—40 Gigabit Ethernet interface

- *fpc*—Flexible PIC Concentrator. EX Series interfaces use the following convention for the FPC number in interface names:
    - On an EX2200 switch, an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, a standalone EX4300 switch, a standalone EX4500, and a standalone EX4550 switch, FPC refers to the switch itself. The FPC number is **0** by default on these switches.
    - On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, an EX4300 Virtual Chassis, an EX4500 Virtual Chassis, an EX4550 Virtual Chassis, or a mixed Virtual Chassis, the FPC number indicates the member ID of the switch in the Virtual Chassis.
    - On an EX6200 switch and a standalone EX8200 switch, the FPC number indicates the slot number of the line card that contains the physical interface. On an EX6200 switch, the FPC number also indicates the slot number of the Switch Fabric and Routing Engine (SRE) module that contains the uplink port.
    - On an EX8200 Virtual Chassis, the FPC number indicates the slot number of the line card on the Virtual Chassis. The line card slots on Virtual Chassis member 0 are numbered 0 through 15; on member 1, they are numbered 16 through 31, and so on.
- *pic*—EX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
    - On EX2200, EX3200, EX3300, EX4200, EX4500 switch, and EX4550 switches, the PIC number is **0** for all built-in interfaces (interfaces that are not uplink ports).
    - On EX2200, EX3200, EX3300, and EX4200 switches, the PIC number is **1** for uplink ports.
    - On EX4300 switches, the PIC number is **0** for built-in network ports, **1** for built-in QSFP+ ports (located on the rear panel of the switch), and **2** for uplink module ports.
    - On EX4500 switches, the PIC number is **1** for ports on the left-hand uplink module and **2** for ports on the right-hand uplink module.
    - On EX4550 switches, the PIC number is **1** for ports in the expansion module or Virtual Chassis module installed in the module slot on the front panel of the switch and **2** for those in the expansion module or Virtual Chassis module installed in the module slot on the rear panel of the switch.
    - On EX6200 and EX8200 switches, the PIC number is always **0**.
- *port*—EX Series interfaces use the following convention for port numbers:
    - On EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, and EX4550 switches, built-in network ports are numbered from left to right. On models that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.
    - Uplink ports in EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, and EX4550 switches are labeled from left to right, starting with **0**.

- On EX6200 and EX8200 switches, the network ports are numbered from left to right on each line card. On line cards that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.
- Uplink ports on an SRE module in an EX6200 switch are labeled from left to right, starting with **0**.

Logical Part of an Interface Name

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

```
Interface    State    VLAN members          Blocking
ge-0/0/0.0   down     remote-analyzer       unblocked
ge-0/0/1.0   down     default               unblocked
ge-0/0/10.0  down     default               unblocked
```

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

- EX Series Switches Interfaces Overview on page 3
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

Understanding Aggregated Ethernet Interfaces and LACP

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

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. The LAG balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, provides additional functionality for LAGs.

This topic describes:

- Link Aggregation Group (LAG) on page 9
- Link Aggregation Control Protocol (LACP) on page 10

Link Aggregation Group (LAG)

You configure a LAG by specifying the link number as a physical device and then associating a set of interfaces (ports) with the link. All the interfaces must have the same speed and be in full-duplex mode. Juniper Networks Junos operating system (Junos OS) for EX Series Ethernet Switches assigns a unique ID and port priority to each interface. The ID and priority are not configurable.

The number of interfaces that can be grouped into a LAG and the total number of LAGs supported on a switch varies according to switch model. Table 5 on page 9 lists the EX Series switches and the maximum number of interfaces per LAG and the maximum number of LAGs they support. MX Series devices can support up to 64 LAGs.

Table 5: Maximum Interfaces per LAG and Maximum LAGs per Switch

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX2200	8	32
EX3200	8	32
EX3300 and EX3300 Virtual Chassis	8	111
EX4200 and EX4200 Virtual Chassis	8	111
EX4300 and EX4300 Virtual Chassis	16	112
EX4500, EX4500 Virtual Chassis, EX4550, and EX4550 Virtual Chassis	8	111
EX6200	8	111
EX8200	12	255
EX8200 Virtual Chassis	12	239

When configuring LAGs, consider the following guidelines:

- You must configure the LAG on both sides of the link.
- You must set the interfaces on either side of the link to the same speed.
- You can configure and apply firewall filters on a LAG.

- You can optionally configure LACP for link negotiation.
- You can optionally configure LACP for link protection.

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a LAG. See *Understanding EX Series Virtual Chassis Port Link Aggregation* and *Understanding Link Aggregation in an EX8200 Virtual Chassis*.

NOTE: **The interfaces that are included within a LAG are sometimes referred to as *member interfaces*. Do not confuse this term with *member switches*, which refers to switches that are interconnected as a Virtual Chassis. It is possible to create a LAG that is composed of member interfaces that are located in different member switches of a Virtual Chassis.**

A LAG hashing algorithm determines how traffic entering a LAG is placed onto the bundle's member links. The LAG hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle. You can configure the fields used by the LAG hashing algorithm on some EX Series switches. See *Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)*.

A LAG creates a single logical point-to-point connection. A typical deployment for a LAG would be to aggregate trunk links between an access switch and a distribution switch or customer edge (CE) router.

Link Aggregation Control Protocol (LACP)

When LACP is configured, it detects misconfigurations on the local end or the remote end of the link. Thus, LACP can help prevent communication failure:

- When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail.
- When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

By default, Ethernet links do not exchange LACP protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit LACP PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when the Ethernet link receives them from the remote end. The transmitting link is known as the *actor* and the receiving link is known as the *partner*.

In a scenario where 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 might 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 PDUs are not received, 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 such a scenario, 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 Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- *Junos OS Network Interfaces Configuration Guide*

Understanding Layer 3 Subinterfaces

A Layer 3 subinterface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 subinterfaces to route traffic among multiple VLANs along a single trunk line that connects a Juniper Networks EX Series Ethernet Switch to a Layer 2 switch. Only one physical connection is required between the switches. This topology is often called a *router on a stick* or a *one-armed router* when the Layer 3 device is a router.

To create Layer 3 subinterfaces on an EX Series switch, you enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

You can partition one physical interface into up to 4094 different subinterfaces, one for each VLAN. We recommend that you use the VLAN ID as the subinterface number when you configure the subinterface. Juniper Networks Junos operating system (Junos OS) reserves VLAN IDs 0 and 4095.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on EX Series switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces. Double-tagging is not supported.

Related Documentation

- EX Series Switches Interfaces Overview on page 3
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 44
- *Junos OS Ethernet Interfaces Configuration Guide*

Understanding Unicast RPF

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. It also helps ensure that traffic arriving on ingress interfaces comes from a network source that the receiving interface can reach.

When you enable unicast RPF, the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF.

NOTE: On Juniper Networks EX3200, EX4200, and EX4300 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see "Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches" on page 15.

This topic covers:

Unicast RPF for Switches Overview

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces.

The type of unicast RPF provided on the switches—that is, strict mode unicast RPF is especially useful on untrusted interfaces. An untrusted interface is an interface where untrusted users or processes can place packets on the network segment.

The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see "When to Enable Unicast RPF" on page 13.)

For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at http://www.ietf.org/rfc/rfc3704.txt.

Unicast RPF Implementation

This section includes:

Unicast RPF Packet Filtering

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

Bootstrap Protocol (BOOTP) and DHCP Requests

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

Default Route Handling

If the best return path to the source is the default route (**0.0.0.0**) and the default route points to **reject**, the switch discards the packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces. A symmetrically routed interface uses the same route in both directions between the source and the destination, as shown in Figure 1 on page 14. Symmetrical routing means that if an interface receives a packet, the switch uses the same interface to send a reply to the packet source (the receiving interface matches the forwarding-table entry for the best return path to the source).

Figure 1: Symmetrically Routed Interfaces

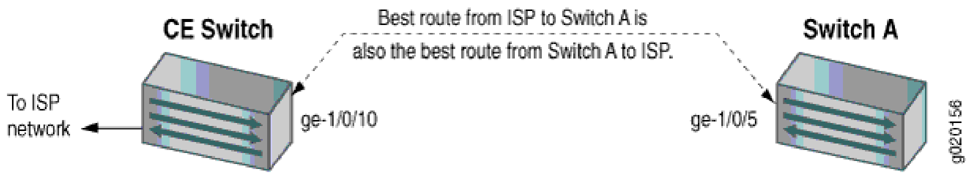

Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link

NOTE: **Because unicast RPF is enabled globally on EX3200, EX4200, and EX4300 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on these switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.**

TIP: **Enabling unicast RPF as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled.**

When Not to Enable Unicast RPF

Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in Figure 2 on page 15. This means that if an interface receives a packet, that interface does not match the forwarding table entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

Figure 2: Asymmetrically Routed Interfaces

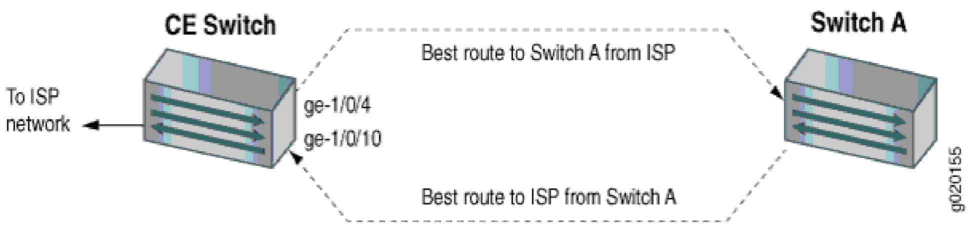

NOTE: Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

Limitations of the Unicast RPF Implementation on EX3200, EX4200, and EX4300 Switches

On EX3200, EX4200, and EX4300 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.

NOTE: You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

QFX switches, OCX switches, and EX3200 and EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths. Using unicast RPF to filter ECMP traffic on these switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

Related Documentation

- Example: Configuring Unicast RPF on an EX Series Switch on page 51
- Configuring Unicast RPF (CLI Procedure) on page 135

- Disabling Unicast RPF (CLI Procedure) on page 137

Understanding IP Directed Broadcast for EX Series Switches

IP directed broadcast helps you implement remote administration tasks such as backups and wake-on-LAN (WOL) application tasks by sending broadcast packets targeted at the hosts in a specified destination subnet. IP directed broadcast packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. When they reach the destination subnet and IP directed broadcast is enabled on the receiving switch, the switch translates (*explodes*) the IP directed broadcast packet into a broadcast that floods the packet on the target subnet. All hosts on the target subnet receive the IP directed broadcast packet.

This topic covers:

IP Directed Broadcast for EX Series Switches Overview

IP directed broadcast packets have a destination IP address that is a valid broadcast address for the subnet that is the target of the directed broadcast (the target subnet). The intent of an IP directed broadcast is to flood the target subnet with the broadcast packets without broadcasting to the entire network. IP directed broadcast packets cannot originate from the target subnet.

When you send an IP directed broadcast packet, as it travels to the target subnet, the network forwards it in the same way as it forwards a unicast packet. When the packet reaches a switch that is directly connected to the target subnet, the switch checks to see whether IP directed broadcast is enabled on the interface that is directly connected to the target subnet:

- If IP directed broadcast is enabled on that interface, the switch broadcasts the packet on that subnet by rewriting the destination IP address as the configured broadcast IP address for the subnet. The switch converts the packet to a link-layer broadcast packet that every host on the network processes.
- If IP directed broadcast is disabled on the interface that is directly connected to the target subnet, the switch drops the packet.

IP Directed Broadcast Implementation for EX Series Switches

You configure IP directed broadcast on a per-subnet basis by enabling IP directed broadcast on the Layer 3 interface of the subnet's VLAN. When the switch that is connected to that subnet receives a packet that has the subnet's broadcast IP address as the destination address, the switch broadcasts the packet to all hosts on the subnet.

By default, IP directed broadcast is disabled.

When to Enable IP Directed Broadcast

IP directed broadcast is disabled by default. Enable IP directed broadcast when you want to perform remote management or administration services such as backups or WOL tasks on hosts in a subnet that does not have a direct connection to the Internet.

Enabling IP directed broadcast on a subnet affects only the hosts within that subnet. Only packets received on the subnet's Layer 3 interface that have the subnet's broadcast IP address as the destination address are flooded on the subnet.

When Not to Enable IP Directed Broadcast

Typically, you do not enable IP directed broadcast on subnets that have direct connections to the Internet. Disabling IP directed broadcast on a subnet's Layer 3 interface affects only that subnet. If you disable IP directed broadcast on a subnet and a packet that has the broadcast IP address of that subnet arrives at the switch, the switch drops the broadcast packet.

If a subnet has a direct connection to the Internet, enabling IP directed broadcast on it increases the network's susceptibility to denial-of-service (DoS) attacks.

For example, a malicious attacker can spoof a source IP address (use a source IP address that is not the actual source of the transmission to deceive a network into identifying the attacker as a legitimate source) and send IP directed broadcasts containing Internet Control Message Protocol (ICMP) echo (ping) packets. When the hosts on the network with IP directed broadcast enabled receive the ICMP echo packets, they all send replies to the victim that has the spoofed source IP address. This creates a flood of ping replies in a DoS attack that can overwhelm the spoofed source address; this is known as a *smurf* attack. Another common DoS attack on exposed networks with IP directed broadcast enabled is a *fraggle* attack, which is similar to a smurf attack except that the malicious packet is a User Datagram Protocol (UDP) echo packet instead of an ICMP echo packet.

Related Documentation

- Example: Configuring IP Directed Broadcast on an EX Series Switch on page 56
- Configuring IP Directed Broadcast (CLI Procedure) on page 137

- *Configuring IP Directed Broadcast (CLI Procedure)*

Understanding Interface Ranges on EX Series Switches

NOTE: This concept uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Understanding Interface Ranges on EX Series Switches*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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 on Juniper Networks EX Series Ethernet switches. 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, 10-Gigabit, and Fast Ethernet 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** node is used in the following configuration hierarchies:

- **ethernet-switching-options analyzer *name* input egress interface**
- **ethernet-switching-options analyzer *name* input ingress interface**
- **ethernet-switching-options analyzer output interface**
- **ethernet-switching-options bpdu-block interface**
- **ethernet-switching-options interfaces**
- **ethernet-switching-options redundant-trunk-group group-name interface**
- **ethernet-switching-options secure-access-port interface**
- **ethernet-switching-options voip interface**
- **poe interface**
- **protocols dot1x authentication interface**
- **protocols gvrp interface**
- **protocols igmp interface**
- **protocols igmp-snooping vlan *vlan-name* interface**
- **protocols isis interface**
- **protocols link-management peer lmp-control-channel interface**
- **protocols link-management te-link *name* interface**
- **protocols lldp interface**
- **protocols lldp-med interface**
- **protocols mpls interface**
- **protocols mstp interface**
- **protocols mstp msti-*id* interface**

- protocols mstp msti-*id* vlan *vlan-id* interface
- protocols oam ethernet link-fault-management interface
- protocols ospf area
- protocols pim interface
- protocols rip group *group-name* neighbor
- protocols ripng group *group-name* neighbor
- protocols router-advertisement interface
- protocols router-discovery interface
- protocols rsvp interface
- protocols sflow interfaces
- protocols stp interface
- protocols vstp vlan *vlan-id* interface
- vlans *vlan-name* interface

Related Documentation

- Interface Ranges on page 104
- EX Series Switches Interfaces Overview on page 3
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 135
- **interface-range on page 217**

802.1Q VLANs Overview

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the Junos OS supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

Related Documentation

- *Configuring Dynamic 802.1Q VLANs*
- *802.1Q VLAN IDs and Ethernet Interface Types*
- *Enabling VLAN Tagging*
- *Binding VLAN IDs to Logical Interfaces*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs*

- *Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface*
- *Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance*
- *Specifying the Interface Over Which VPN Traffic Travels to the CE Router*
- *Specifying the Interface to Handle Traffic for a CCC*
- *Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface*
- *Configuring a VLAN-Bundled Logical Interface to Support a Layer 2 VPN Routing Instance*
- *Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit*
- *Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface*
- *Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface*
- *Configuring a Logical Interface for Access Mode*
- *Configuring a Logical Interface for Trunk Mode*
- *Configuring the VLAN ID List for a Trunk Interface*
- *Configuring a Trunk Interface on a Bridge Network*
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 20
- GRE Tunneling on page 21
- Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches on page 23
- Configuration Limitations on page 23

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 23.)

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 21
- Number of Source and Destination Tunnels Allowed on a Switch on page 21
- Class of Service on GRE Tunnels on page 22
- Applying Firewall Filters to GRE Traffic on page 22

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 6 on page 22 identifies these constraints.

Table 6: 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 7 on page 23 lists features that are not supported with GRE.

Table 7: 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) on page 142
- *Configuring Generic Routing Encapsulation Tunneling*
- *Configuring a Firewall Filter to De-Encapsulate GRE Traffic on a QFX5100 or OCX Switch*

Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches

Streaming video technology was introduced in 1997. Multicast protocols were subsequently developed to reduce data replication and network overloads. With multicasting, servers can send a single stream to a group of recipients instead of sending multiple unicast streams. While the use of streaming video technology was previously limited to occasional company presentations, multicasting has provided a boost to the technology resulting in a constant stream of movies, real-time data, news clips, and amateur videos flowing nonstop to computers, TVs, tablets, and phones. However, all of these streams quickly overwhelmed the capacity of network hardware and increased bandwidth demands leading to unacceptable blips and stutters in transmission.

To satisfy the growing bandwidth demands, multiple links were virtually aggregated to form bigger logical point-to-point link channels for the flow of data. These virtual link combinations are called multicast interfaces, also known as link aggregation groups (LAGs).

Multicast load balancing involves managing the individual links in each LAG to ensure that each link is used efficiently. Hashing algorithms continually evaluate the data stream, adjusting stream distribution over the links in the LAG, so that no link is underutilized or overutilized. Multicast load balancing is enabled by default on Juniper Networks EX8200 Ethernet Switches.

This topic includes:

Create LAGs for Multicasting in Increments of 10 Gigabits

The maximum link size on an EX8200 switch is 10 gigabits. If you need a larger link on an EX8200 switch, you can combine up to twelve 10-gigabit links. In the sample topology shown in Figure 3 on page 25, four 10-gigabit links have been aggregated to form each 40-gigabit link.

Figure 3: 40-Gigabit LAGs on EX8200 Switches

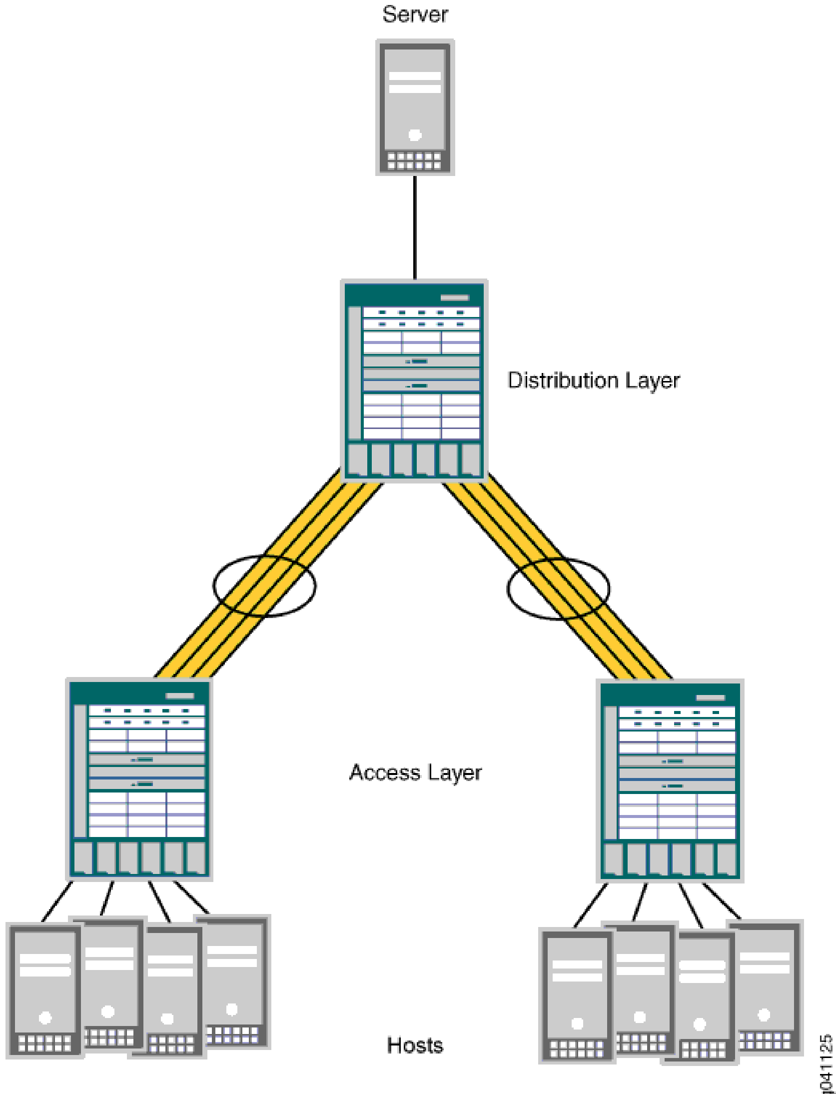

When Should I Use Multicast Load Balancing?

Use a LAG with multicast load balancing when you need a downstream link greater than 10 gigabits. This need frequently arises when you act as a service provider or when you multicast video to a large audience.

To use multicast load balancing, you need the following:

- An EX8200 switch—Standalone switches support multicast load balancing, while Virtual Chassis does not.
- A Layer 3 routed multicast setup—For information about configuring multicasting, see *Junos OS Routing Protocols Configuration Guide*.

- Aggregated 10-gigabit links in a LAG—For information about configuring LAGs with multicast load balancing , see "Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)" on page 144.

How Does Multicast Load Balancing Work?

Juniper Networks Junos operating system (Junos OS) supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs and is supported only on Layer 3 interfaces. When traffic can use multiple member links, traffic that is part of the same stream must always be on the same link.

Multicast load balancing uses one of seven available hashing algorithms and a technique called queue shuffling (alternating between two queues) to distribute and balance the data, directing streams over all available aggregated links. You can select one of the seven algorithms when you configure multicast load balancing, or you can use the default algorithm, crc-sgip, which uses a cyclic redundancy check (CRC) algorithm on the multicast packets' group IP address. We recommend that you start with the crc-sgip default and try other options if this algorithm does not evenly distribute the Layer 3 routed multicast traffic. Six of the algorithms are based on the hashed value of IP addresses (IPv4 or IPv6) and will produce the same result each time they are used. Only the balanced mode option produces results that vary depending on the order in which streams are added. See Table 8 on page 26 for more information.

Table 8: Hashing Algorithms Used by Multicast Load Balancing

Hashing Algorithms	Based On	Best Use
crc-sgip	Cyclic redundancy check of multicast packets' source and group IP address	Default—high-performance management of IP traffic on 10-Gigabit Ethernet network. Predictable assignment to the same link each time. This mode is complex but yields a good distributed hash.
crc-gip	Cyclic redundancy check of multicast packets' group IP address	Predictable assignment to the same link each time. Try this mode when crc-sgip does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
crc-sip	Cyclic redundancy check of multicast packets' source IP address	Predictable assignment to the same link each time. Try this mode when crc-sgip does not evenly distribute the Layer 3 routed multicast traffic and the stream sources vary.
simple-sgip	XOR calculation of multicast packets' source and group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as crc-sgip yields. Try this mode when crc-sgip does not evenly distribute the Layer 3 routed multicast traffic.
simple-gip	XOR calculation of multicast packets' group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as crc-gip yields. Try this when crc-gip does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.

Table 8: Hashing Algorithms Used by Multicast Load Balancing *(continued)*

Hashing Algorithms	Based On	Best Use
simple-sip	XOR calculation of multicast packets' source IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as crc-sip yields. Try this mode when crc-sip does not evenly distribute the Layer 3 routed multicast traffic and stream sources vary.
balanced	Round-robin calculation method used to identify multicast links with the least amount of traffic	Best balance is achieved, but you cannot predict which link will be consistently used because that depends on the order in which streams come online. Use when consistent assignment is not needed after every reboot.

How Do I Implement Multicast Load Balancing on an EX8200 Switch?

To implement multicast load balancing with an optimized level of throughput on an EX8200 switch, follow these recommendations:

- Allow 25 percent unused bandwidth in the aggregated link to accommodate any dynamic imbalances due to link changes caused by sharing multicast interfaces.
- For downstream links, use multicast interfaces of the same size whenever possible. Also, for downstream aggregated links, throughput is optimized when members of the aggregated link belong to the same devices.
- For upstream aggregated links, use a Layer 3 link whenever possible. Also, for upstream aggregated links, throughput is optimized when the members of the aggregated link belong to different devices.

Related Documentation

- Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches on page 59
- Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure) on page 144

Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces

Energy Efficient Ethernet (EEE), an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, reduces the power consumption of physical layer devices (PHYs) during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when the link is idle.

An Ethernet link consumes power even when a ink is idle. EEE provides a method to utilize power in such a way that Ethernet links use power only during data transmission. EEE specifies a signaling protocol, Low Power Idle (LPI) for achieving the power saving during the idle time of Ethernet links. EEE allows PHYs to exchange LPI indications to signal the transition to low power mode when there is no traffic. LPI indicates when a link can go idle and when the link needs to resume after a predefined delay without impacting data transmission.

The following copper PHYs are standardized by IEEE 802.3az:

- 100BASE-T
- 1000BASE-T
- 10GBASE-T

Related Documentation

- Configuring Energy Efficient Ethernet on Interfaces (CLI Procedure) on page 145

Understanding Local Link Bias

Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a link aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF using a different member link in the LAG bundle. The flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is enabled is illustrated in Figure 4 on page 28.

Figure 4: Egress Traffic Flow with Local Link Bias

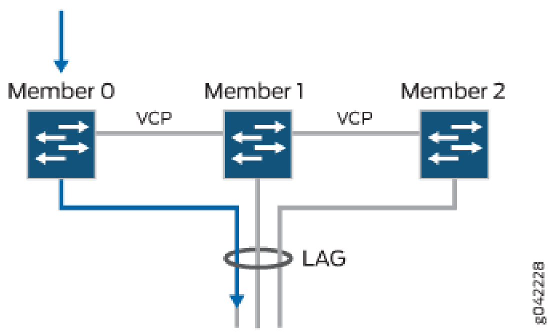

When local link bias is disabled, egress traffic exiting a Virtual Chassis or VCF on a LAG bundle can be forwarded out of any member link in the LAG bundle. Traffic forwarding decisions are made by an internal algorithm that attempts to load-balance traffic between the member links in the bundle. VCP bandwidth is frequently consumed by egress traffic when local link bias is disabled because the egress traffic traverses the VCPs to reach the destination egress member link in the LAG bundle. The flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is disabled is illustrated in Figure 5 on page 29.

Figure 5: Egress Traffic Flow without Local Link Bias

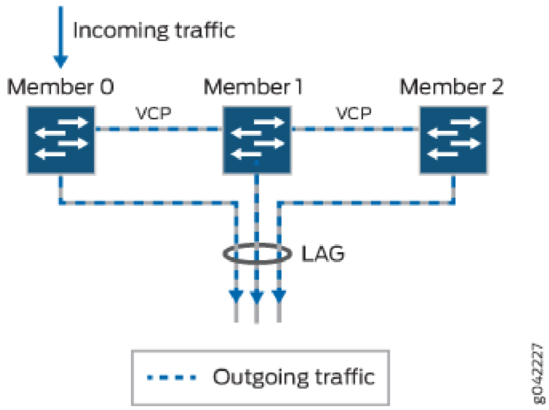

Starting in Junos OS Release 14.1X53-D25, local link bias can be enabled globally for all LAG bundles in a Virtual Chassis or VCF, or individually per LAG bundle in a Virtual Chassis. In prior Junos OS releases, local link bias could be enabled individually per LAG bundle only.

A Virtual Chassis or VCF that has multiple LAG bundles can contain bundles that have and have not enabled local link bias. Local link bias only impacts the forwarding of unicast traffic exiting a Virtual Chassis or VCF; ingress traffic handling is not impacted by the local link bias setting. Egress multicast, unknown unicast, and broadcast traffic exiting a Virtual Chassis or VCF over a LAG bundle is not impacted by the local link bias setting and is always load-balanced among the member links. Local link bias is disabled, by default.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on the LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced across the member links in the LAG bundle as it exits the Virtual Chassis or VCF.

Related Documentation

- Configuring Local Link Bias (CLI Procedure) on page 147

PART 2

Configuration

CHAPTER 2

Configuration Examples

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your EX Series switch model. See "Understanding Aggregated Ethernet Interfaces and LACP" on page 8 for more information.

This example describes how to configure uplink LAGs to connect a Virtual Chassis access switch to a Virtual Chassis distribution switch:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches
- Two EX4200-48P switches
- Two EX4200-24F switches
- Four XFP uplink modules

Before you configure the LAGs, be sure you have:

- Configured the Virtual Chassis switches. See *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*.
- Configured the uplink ports on the switches as trunk ports. See "Configuring Gigabit Ethernet Interfaces (CLI Procedure)" on page 66.

Overview and Topology

For maximum speed and resiliency, you can combine uplinks between an access switch and a distribution switch into LAGs. Using LAGs can be particularly effective when connecting a multimember Virtual Chassis access switch to a multimember Virtual Chassis distribution switch.

The Virtual Chassis access switch in this example is composed of two member switches. Each member switch has an uplink module with two 10-Gigabit Ethernet ports. These ports are configured as trunk ports, connecting the access switch with the distribution switch.

Configuring the uplinks as LAGs has the following advantages:

- Link Aggregation Control Protocol (LACP) can optionally be configured for link negotiation.
- It doubles the speed of each uplink from 10 Gbps to 20 Gbps.
- If one physical port is lost for any reason (a cable is unplugged or a switch port fails, or one member switch is unavailable), the logical port transparently continues to function over the remaining physical port.

The topology used in this example consists of one Virtual Chassis access switch and one Virtual Chassis distribution switch. The access switch is composed of two EX4200-48P switches (SWA-0 and SWA-1), interconnected to each other with their Virtual Chassis ports (VCPs) as member switches of Host-A. The distribution switch is composed of two EX4200-24F switches (SWD-0 and SWD-1), interconnected with their VCPs as member switches of Host-D.

Each member of the access switch has an uplink module installed. Each uplink module has two ports. The uplinks are configured to act as trunk ports, connecting the access switch with the distribution switch. One uplink port from SWA-0 and one uplink port from SWA-1 are combined as LAG **ae0** to SWD-0. This link is used for one VLAN. The remaining uplink ports from SWA-0 and from SWA-1 are combined as a second LAG connection (**ae1**) to SWD-1. LAG **ae1** is used for another VLAN.

NOTE: 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.

Figure 6: Topology for LAGs Connecting an EX4200 Virtual Chassis Access Switch to an EX4200 Virtual Chassis Distribution Switch

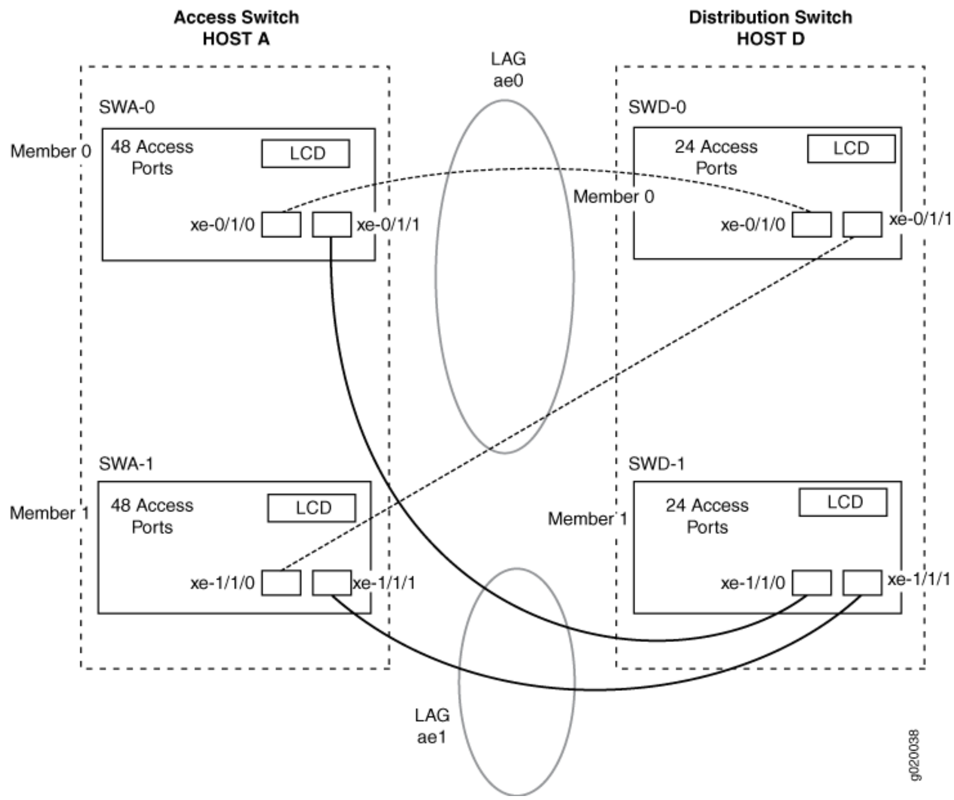

Table 9 on page 35 details the topology used in this configuration example.

Table 9: Components of the Topology for Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWA-0	Host-A Access switch VCID 1	EX4200-48P switch	One XFP uplink module	0	**xe-0/1/0** to SWD-0 **xe-0/1/1** to SWD-1
SWA-1	Host-A Access switch VCID 1	EX4200-48P switch	One XFP uplink module	1	**xe-1/1/0** to SWD-0 **xe-1/1/1** to SWD-1

Table 9: Components of the Topology for Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch *(continued)*

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWD-0	Host-D Distribution switch VCID 4	EX4200 L-24F switch	One XFP uplink module	0	**xe-0/1/0** to SWA-0 **xe-0/1/1** to SWA-1
SWD-1	Host-D Distribution switch VCID 4	EX4200 L-24F switch	One XFP uplink module	1	**xe-1/1/0** to SWA-0 **xe-1/1/1** to SWA-1

Configuration

To configure two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch:

CLI Quick Configuration

To quickly configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis aggregated-devices ethernet device-count 2
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options minimum-links 1
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family inet address 192.0.2.0/25
set interfaces ae1 unit 0 family inet address 192.0.2.128/25
set interfaces xe-0/1/0 ether-options 802.3ad ae0
set interfaces xe-1/1/0 ether-options 802.3ad ae0
set interfaces xe-0/1/1 ether-options 802.3ad ae1
set interfaces xe-1/1/1 ether-options 802.3ad ae1
```

Step-by-Step Procedure

To configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch:

1. Specify the number of LAGs to be created on the chassis:

   ```
   [edit chassis]
   user@Host-A# set aggregated-devices ethernet device-count 2
   ```
2. Specify the number of links that need to be present for the **ae0** LAG interface to be **up**:

   ```
   [edit interfaces]
   user@Host-A# set ae0 aggregated-ether-options minimum-links 1
   ```
3. Specify the number of links that need to be present for the **ae1** LAG interface to be **up**:

   ```
   [edit interfaces]
   user@Host-A# set ae1 aggregated-ether-options minimum-links 1
   ```
4. Specify the media speed of the **ae0** link:

   ```
   [edit interfaces]
   user@Host-A# set ae0 aggregated-ether-options link-speed 10g
   ```


5. Specify the media speed of the **ae1** link:

   ```
   [edit interfaces]
   user@Host-A# set ae1 aggregated-ether-options link-speed 10g
   ```

6. Specify the interface ID of the uplinks to be included in LAG **ae0**:

   ```
   [edit interfaces]
   user@Host-A# set xe-0/1/0 ether-options 802.3ad ae0
   user@Host-A# set xe-1/1/0 ether-options 802.3ad ae0
   ```

7. Specify the interface ID of the uplinks to be included in LAG **ae1**:

   ```
   [edit interfaces]
   user@Host-A# set xe-0/1/1 ether-options 802.3ad ae1
   user@Host-A# set xe-1/1/1 ether-options 802.3ad ae1
   ```

8. Specify that LAG **ae0** belongs to the subnet for the employee broadcast domain:

   ```
   [edit interfaces]
   user@Host-A# set ae0 unit 0 family inet address 192.0.2.0/25
   ```

9. Specify that LAG **ae1** belongs to the subnet for the guest broadcast domain:

   ```
   [edit interfaces]
   user@Host-A# set ae1 unit 0 family inet address 192.0.2.128/25
   ```

Results

Display the results of the configuration:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  ae0 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
      family inet {
        address 192.0.2.0/25;
      }
    }
  }
  ae1 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 1;
    }
    unit 0 {
      family inet {
        address 192.0.2.128/25;
      }
    }
  xe–0/1/0 {
```

```
        ether-options {
          802.3ad ae0;
        }
      }
      xe–1/1/0 {
        ether-options {
          802.3ad ae0;
        }
      }
      xe–0/1/1 {
        ether-options {
          802.3ad ae1;
        }
      }
      xe–1/1/1 {
        ether-options {
          802.3ad ae1;
        }
      }
    }
```

Verification

To verify that switching is operational and two LAGs have been created, perform these tasks:

- Verifying That LAG ae0 Has Been Created on page 38
- Verifying That LAG ae1 Has Been Created on page 38

Verifying That LAG ae0 Has Been Created

Purpose Verify that LAG **ae0** has been created on the switch.

Action **show interfaces ae0 terse**

```
Interface              Admin      Link  Proto      Local              Remote
ae0                    up         up
ae0.0                  up         up      inet     192.0.2.0/25
```

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

Verifying That LAG ae1 Has Been Created

Purpose Verify that LAG **ae1** has been created on the switch

Action **show interfaces ae1 terse**

```
Interface                Admin Link Proto     Local                 Remote
ae1                      up    down
ae1.0                    up    down inet    192.0.2.128/25
```

Meaning The output shows that the **ae1** link is down.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem 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).
- 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 (or the same Virtual Chassis).

Related Documentation

- *Example: Configuring an EX4200 Virtual Chassis with a Master and Backup in a Single Wiring Closet*
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- *Example: Connecting an Access Switch to a Distribution Switch*.
- *Virtual Chassis Cabling Configuration Examples for EX4200 Switches*
- *Installing an Uplink Module in an EX4200 Switch*
- *Uplink Modules in EX4200 Switches*

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch

EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. EX Series switches allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in "Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch" on page 33:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.0 or later for EX Series switches
- Two EX4200-48P switches
- Two EX4200-24F switches
- Four EX Series XFP uplink modules

Before you configure LACP, be sure you have:

- Set up the Virtual Chassis switches. See *Configuring an EX4200, EX4500, or EX4550 Virtual Chassis (CLI Procedure)*.
- Configured the uplink ports on the switches as trunk ports. See "Configuring Gigabit Ethernet Interfaces (CLI Procedure)" on page 66.
- Configured the LAGs. See "Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch" on page 33.

Overview and Topology

This example assumes that you are familiar with "Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch" on page 33. The topology in this example is exactly the same as the topology in that other example. This example shows how to use LACP to enhance the LAG functionality.

LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP mode can be either active or passive.

NOTE: **If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and responses to LACP packets, you must enable LACP in active mode.**

By default, the actor and partner send LACP packets every second.

The interval can be fast (every second) or slow (every 30 seconds).

Configuring LACP for the LAGs on the Virtual Chassis Access Switch

To configure LACP for the access switch LAGs, perform these tasks:

CLI Quick Configuration

To quickly configure LACP for the access switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
set interfaces ae1 aggregated-ether-options lacp active periodic fast
```

Step-by-Step Procedure

To configure LACP for Host-A LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-A#set ae0 aggregated-ether-options lacp active periodic fast
user@Host-A#set ae1 aggregated-ether-options lacp active periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
user@Host-A# show
ae0 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
```

Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch

To configure LACP for the two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch, perform these tasks:

CLI Quick Configuration

To quickly configure LACP for the distribution switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit interfaces]
set ae0 aggregated-ether-options lacp passive periodic fast
set ae1 aggregated-ether-options lacp passive periodic fast
```

Step-by-Step Procedure

To configure LACP for Host D LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-D#set ae0 aggregated-ether-options lacp passive periodic fast
user@Host-D#set ae1 aggregated-ether-options lacp passive periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
```

```
user@Host-D# show
ae0 {
  aggregated-ether-options {
    lacp {
      passive;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      passive
      periodic fast;
    }
  }
}
```

Verification

To verify that LACP packets are being exchanged, perform these tasks:

- Verifying the LACP Settings on page 42
- Verifying That the LACP Packets Are Being Exchanged on page 42

Verifying the LACP Settings

Purpose Verify that LACP has been set up correctly.

Action Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@Host-A> show lacp interfaces xe-0/1/0

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 The output indicates that LACP has been set up correctly and is active at one end.

Verifying That the LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged.

Action Use the **show interfaces ae*x* statistics** command to display LACP information.

```
user@Host-A> show interfaces ae0 statistics

Physical interface: ae0, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 30
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
    Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no protocol data units (PDUs) are being exchanged.

Troubleshooting

To troubleshoot a nonworking LACP link, perform these tasks:

- Troubleshooting a Nonworking LACP Link on page 43

Troubleshooting a Nonworking LACP Link

Problem The LACP link is not working.

Solution Check the following:

- Remove the LACP configuration and verify whether the static LAG is up.
- Verify that LACP is configured at both ends.
- Verify that LACP is not passive at both ends.
- Verify whether LACP protocol data units (PDUs) are being exchanged by running the **monitor traffic-interface *lag-member* detail** command.

Related Documentation

- *Example: Connecting an Access Switch to a Distribution Switch*
- *Virtual Chassis Cabling Configuration Examples for EX4200 Switches*

- *Installing an Uplink Module in an EX4200 Switch*
- Understanding Aggregated Ethernet Interfaces and LACP on page 8

Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch

In a large LAN, you commonly need to partition the network into multiple VLANs. You can configure Layer 3 subinterfaces to route traffic between the VLANs. In one common topology, known as a "router on a stick" or a "one-armed router," you connect a router to an access switch with connections to multiple VLANs.

This example describes how to create Layer 3 subinterfaces on trunk interfaces of a distribution switch and access switch so that you can route traffic among multiple VLANs:

Requirements

This example uses the following hardware and software components:

- For the distribution switch, one EX4200-24F switch. This model is designed to be used as a distribution switch for aggregation or collapsed core network topologies and in space-constrained data centers. It has twenty-four 1-Gigabit Ethernet fiber SFP ports and an EX-UM-2XFP uplink module with two 10-Gigabit Ethernet XFP ports.
- For the access switch, any Layer 2 switch that supports 802.1Q VLAN tags.
- Junos OS Release 9.2 or later for EX Series switches.

Before you connect the switches, make sure you have:

- Connected the two switches.
- Configured the necessary VLANs. See *Configuring VLANs for EX Series Switches (CLI Procedure)* or *Configuring VLANs for EX Series Switches (J-Web Procedure)*.

Overview and Topology

In a large office with multiple buildings and VLANs, you commonly aggregate traffic from a number of access switches into a distribution switch. This configuration example shows a simple topology to illustrate how to connect a single Layer 2 access switch connected to multiple VLANs to a distribution switch, enabling traffic to pass between those VLANs.

In the example topology, the LAN is segmented into five VLANs, all associated with interfaces on the access switch. One 1-Gigabit Ethernet port on the access switch's uplink module connects to one 1-Gigabit Ethernet port on the distribution switch.

Table 10 on page 45 lists the settings for the example topology.

Table 10: Components of the Topology for Creating Layer 3 Subinterfaces on an Access Switch and a Distribution Switch

Property	Settings
Access switch hardware	Any Layer 2 switch with multiple 1-Gigabit Ethernet ports and at least one 1-Gigabit Ethernet uplink module
Distribution switch hardware	EX4200-24F, 24 1-Gigabit Ethernet fiber SPF ports (ge-0/0/0 through ge-0/0/23); one 2-port 10-Gigabit Ethernet XFP uplink module (EX-UM-4SFP)
VLAN names and tag IDs	vlan1, tag 101 vlan2, tag 102 vlan3, tag 103 vlan4, tag 104 vlan5, tag 105
VLAN subnets	vlan1: 1.1.1.0/24 (addresses 1.1.1.1 through 1.1.1.254) vlan2: 2.1.1.0/24 (addresses 2.1.1.1 through 2.1.1.254) vlan3: 3.1.1.0/24 (addresses 3.1.1.1 through 3.1.1.254) vlan4: 4.1.1.0/24 (addresses 4.1.1.1 through 4.1.1.254) vlan5: 5.1.1.0/24 (addresses 5.1.1.1 through 5.1.1.254)
Port interfaces	On the access switch: ge-0/1/0 On the distribution switch: ge-0/0/0

Configuring the Access Switch Subinterfaces

CLI Quick Configuration

To quickly create and configure subinterfaces on the access switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/1/0 vlan-tagging
set interfaces ge-0/1/0 unit 0 vlan-id 101 family inet address 1.1.1.1/24
set interfaces ge-0/1/0 unit 1 vlan-id 102 family inet address 2.1.1.1/24
set interfaces ge-0/1/0 unit 2 vlan-id 103 family inet address 3.1.1.1/24
set interfaces ge-0/1/0 unit 3 vlan-id 104 family inet address 4.1.1.1/24
set interfaces ge-0/1/0 unit 4 vlan-id 105 family inet address 5.1.1.1/24
```

Step-by-Step Procedure

To configure the subinterfaces on the access switch:

1. On the trunk interface of the access switch, enable VLAN tagging:

   ```
   [edit interfaces ge-0/1/0]
   user@access-switch# set vlan-tagging
   ```
2. Bind vlan1's VLAN ID to the logical interface:

   ```
   [edit interfaces ge-0/1/0]
   user@access-switch# set unit 0 vlan-id 101
   ```

3. Set vlan1's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 0 family inet address 1.1.1.1/24
```

4. Bind vlan2's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 1 vlan-id 102
```

5. Set vlan2's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 1 family inet address 2.1.1.1/24
```

6. Bind vlan3's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 vlan–id 103
```

7. Set vlan3's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 family inet address 3.1.1.1/24
```

8. Bind vlan4's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 vlan-id 104
```

9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 family inet address 4.1.1.1/24
```

10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 4 vlan-id 105
```

11. Set vlan5's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 4 family inet address 5.1.1.1/24
```

Results

Check the results of the configuration:

```
user@access-switch> show configuration
interfaces {
  ge-0/1/0 {
    vlan-tagging;
    unit 0 {
      vlan-id 101;
      family inet {
        address 1.1.1.1/24;
      }
    }
    unit 1 {
      vlan-id 102;
      family inet {
        address 2.1.1.1/24;
      }
    }
    unit 2 {
      vlan-id 103;
      family inet {
```

```
            address 3.1.1.1/24;
        }
    }
    unit 3 {
        vlan-id 104;
        family inet {
            address 4.1.1.1/24;
        }
    }
    unit 4 {
        vlan-id 105;
        family inet {
            address 5.1.1.1/24;
        }
    }
}
```

Configuring the Distribution Switch Subinterfaces

CLI Quick Configuration

To quickly create and configure subinterfaces on the distribution switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 vlan-id 101 family inet address 1.1.1.2/24
set interfaces ge-0/0/0 unit 1 vlan-id 102 family inet address 2.1.1.2/24
set interfaces ge-0/0/0 unit 2 vlan-id 103 family inet address 3.1.1.2/24
set interfaces ge-0/0/0 unit 3 vlan-id 104 family inet address 4.1.1.2/24
set interfaces ge-0/0/0 unit 4 vlan-id 105 family inet address 5.1.1.2/24
```

Step-by-Step Procedure

To configure subinterfaces on the distribution switch:

1. On the trunk interface of the distribution switch, enable VLAN tagging:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set vlan-tagging
```

2. Bind vlan1's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 0 vlan-id 101
```

3. Set vlan1's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 0 family inet address 1.1.1.2/24
```

4. Bind vlan2's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 1 vlan-id 102
```

5. Set vlan2's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 1 family inet address 2.1.1.2/24
```

6. Bind vlan3's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 2 vlan-id 103
```

7. Set vlan3's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 2 family inet address 3.1.1.2/24
```

8. Bind vlan4's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 3 vlan-id 104
```

9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 3 family inet address 4.1.1.2/24
```

10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 4 vlan-id 105
```

11. Set vlan5's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 4 family inet address 5.1.1.2/24
```

Results

```
user@distribution-switch> show configuration
interfaces {
  ge-0/0/0 {
    vlan-tagging;
    unit 0 {
      vlan-id 101;
      family inet {
        address 1.1.1.2/24;
      }
    }
    unit 1 {
      vlan-id 102;
      family inet {
        address 2.1.1.2/24;
      }
    }
    unit 2 {
      vlan-id 103;
      family inet {
        address 3.1.1.2/24;
      }
    }
    unit 3 {
      vlan-id 104;
      family inet {
        address 4.1.1.2/24;
      }
    }
    unit 4 {
      vlan-id 105;
      family inet {
        address 5.1.1.2/24;
      }
    }
  }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Subinterfaces Were Created on page 49
- Verifying That Traffic Passes Between VLANs on page 49

Verifying That Subinterfaces Were Created

Purpose Verify that the subinterfaces were properly created on the access switch and distribution switch.

Action

1. Use the **show interfaces** command on the access switch:

```
user@access-switch> show interfaces ge-0/1/0 terse
Interface               Admin Link Proto    Local                 Remote
ge-0/1/0                up    up
ge-0/1/0.0              up    up   inet     1.1.1.1/24
ge-0/1/0.1              up    up   inet     2.1.1.1/24
ge-0/1/0.2              up    up   inet     3.1.1.1/24
ge-0/1/0.3              up    up   inet     4.1.1.1/24
ge-0/1/0.4              up    up   inet     5.1.1.1/24
ge-0/1/0.32767          up    up
```

2. Use the show interfaces command on the distribution switch:

```
user@distribution-switch> show interfaces ge-0/0/0 terse
Interface               Admin Link Proto    Local                 Remote
ge-0/0/0                up    up
ge-0/0/0.0              up    up   inet     1.1.1.2/24
ge-0/0/0.1              up    up   inet     2.1.1.2/24
ge-0/0/0.2              up    up   inet     3.1.1.2/24
ge-0/0/0.3              up    up   inet     4.1.1.2/24
ge-0/0/0.4              up    up   inet     5.1.1.2/24
ge-0/0/0.32767          up    up
```

Meaning Each subinterface created is displayed as a **ge-*fpc/pic/port.x*** logical interface, where **x** is the unit number in the configuration. The status is listed as **up**, indicating the link is working.

Verifying That Traffic Passes Between VLANs

Purpose Verify that the distribution switch is correctly routing traffic from one VLAN to another.

Action Ping from the access switch to the distribution switch on each subinterface.

1. From the access switch, ping the address of the vlan1 subinterface on the distribution switch:

```
user@access-switch> ping 1.1.1.2 count 4
PING 1.1.1.2 (1.1.1.2): 56 data bytes
64 bytes from 1.1.1.2: icmp_seq=0 ttl=64 time=0.333 ms
64 bytes from 1.1.1.2: icmp_seq=1 ttl=64 time=0.113 ms
64 bytes from 1.1.1.2: icmp_seq=2 ttl=64 time=0.112 ms
64 bytes from 1.1.1.2: icmp_seq=3 ttl=64 time=0.158 ms

--- 1.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.112/0.179/0.333/0.091 ms
```

2. From the access switch, ping the address of the vlan2 subinterface on the distribution switch:

```
user@access-switch> ping 2.1.1.2 count 4
PING 2.1.1.2 (2.1.1.2): 56 data bytes
64 bytes from 2.1.1.2: icmp_seq=0 ttl=64 time=0.241 ms
64 bytes from 2.1.1.2: icmp_seq=1 ttl=64 time=0.113 ms
64 bytes from 2.1.1.2: icmp_seq=2 ttl=64 time=0.162 ms
64 bytes from 2.1.1.2: icmp_seq=3 ttl=64 time=0.167 ms

--- 2.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.113/0.171/0.241/0.046 ms
```

3. From the access switch, ping the address of the vlan3 subinterface on the distribution switch:

```
user@access-switch> ping 3.1.1.2 count 4
PING 3.1.1.2 (3.1.1.2): 56 data bytes
64 bytes from 3.1.1.2: icmp_seq=0 ttl=64 time=0.341 ms
64 bytes from 3.1.1.2: icmp_seq=1 ttl=64 time=0.162 ms
64 bytes from 3.1.1.2: icmp_seq=2 ttl=64 time=0.112 ms
64 bytes from 3.1.1.2: icmp_seq=3 ttl=64 time=0.208 ms

--- 3.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.112/0.206/0.341/0.085 ms
```

4. From the access switch, ping the address of the vlan4 subinterface on the distribution switch:

```
user@access-switch> ping 4.1.1.2 count 4
PING 4.1.1.2 (4.1.1.2): 56 data bytes
64 bytes from 4.1.1.2: icmp_seq=0 ttl=64 time=0.226 ms
64 bytes from 4.1.1.2: icmp_seq=1 ttl=64 time=0.166 ms
64 bytes from 4.1.1.2: icmp_seq=2 ttl=64 time=0.107 ms
64 bytes from 4.1.1.2: icmp_seq=3 ttl=64 time=0.221 ms

--- 4.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.107/0.180/0.226/0.048 ms
```

5. From the access switch, ping the address of the vlan5 subinterface on the distribution switch:

```
user@access-switch> ping 5.1.1.2 count 4
```



```
PING 5.1.1.2 (5.1.1.2): 56 data bytes
64 bytes from 5.1.1.2: icmp_seq=0 ttl=64 time=0.224 ms
64 bytes from 5.1.1.2: icmp_seq=1 ttl=64 time=0.104 ms
64 bytes from 5.1.1.2: icmp_seq=2 ttl=64 time=0.102 ms
64 bytes from 5.1.1.2: icmp_seq=3 ttl=64 time=0.170 ms

--- 5.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.102/0.150/0.224/0.051 ms
```

Meaning If all the ping packets are transmitted and are received by the destination address, the subinterfaces are up and working.

Related Documentation

- *Understanding Layer 3 Logical Interfaces*
- *Example: Connecting an Access Switch to a Distribution Switch*
- *Configuring a Layer 3 Logical Interface*
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 135

Example: Configuring Unicast RPF on an EX Series Switch

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend the switch ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast reverse-path forwarding (RPF) on a customer-edge interface to filter incoming traffic:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 10.1 or later for EX Series switches
- Two EX8200 switches

Before you begin, be sure you have:

- Connected the two switches by symmetrically routed interfaces.
- Ensured that the interface on which you will configure unicast RPF is symmetrically routed.

Overview and Topology

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a denial-of-service (DoS) attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and distributed denial-of-service (DDoS) attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding-table entry for its source address. If the switch uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the switch forwards the packet. If the switch does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the switch discards the packet.

This example uses two EX8200 switches. On EX3200 and EX4200 switches, you cannot configure individual interfaces for unicast RPF. On EX3200 and EX4200 switches, the switch applies unicast RPF globally to all interfaces on the switch. See "Understanding Unicast RPF" on page 12 for more information on limitations regarding the configuration of unicast RPF on EX3200 and EX4200 switches.

In this example, an enterprise network's system administrator wants to protect Switch A against potential DoS and DDoS attacks from the Internet. The administrator configures unicast RPF on interface **ge-1/0/10** on Switch A. Packets arriving on interface **ge-1/0/10** on Switch A from the Switch B source also use incoming interface **ge-1/0/10** as the best return path to send packets back to the source.

The topology of this configuration example uses two EX8200 switches, Switch A and Switch B, connected by symmetrically routed interfaces:

- Switch A is on the edge of an enterprise network. The interface **ge-1/0/10** on Switch A connects to the interface **ge-1/0/5** on Switch B.
- Switch B is on the edge of the service provider network that connects the enterprise network to the Internet.

Configuration

To enable unicast RPF, perform these tasks:

CLI Quick Configuration

To quickly configure unicast RPF on Switch A, copy the following command and paste it into the switch terminal window:

```
[edit interfaces]
set ge-1/0/10 unit 0 family inet rpf-check
```

Step-by-Step Procedure

To configure unicast RPF on Switch A:

1. Enable unicast RPF on interface **ge-1/0/10**:

```
[edit interfaces]
user@switch# set ge-1/0/10 unit 0 family inet rpf-check
```

Results

Check the results:

```
[edit interfaces]
user@switch# show
ge-1/0/10 {
   unit 0 {
      family inet {
         rpf-check;
      }
   }
}
```

Verification

To confirm that the configuration is correct, perform these tasks:

- Verifying That Unicast RPF Is Enabled on the Switch on page 53

Verifying That Unicast RPF Is Enabled on the Switch

Purpose

Verify that unicast RPF is enabled.

Action Verify that unicast RPF is enabled on interface **ge-1/0/10** by using the **show interfaces ge-1/0/10 extensive** or **show interfaces ge-1/0/10 detail** command.

```
user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
  Interface index: 139, SNMP ifIndex: 58, Generation: 140
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  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: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
  Last flapped   : Never
  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, 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  : 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
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: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit 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
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
      Protocol inet, Generation: 144, Route table: 0
    Flags: uRPF
    Addresses, Flags: Is-Preferred Is-Primary
```

Meaning The second-to-last line of the display shows the unicast RPF flag enabled, confirming that unicast RPF is enabled on interface **ge-1/0/10**.

Related Documentation

- Configuring Unicast RPF (CLI Procedure) on page 135
- Disabling Unicast RPF (CLI Procedure) on page 137

Example: Configuring IP Directed Broadcast on an EX Series Switch

IP directed broadcast provides a method of sending broadcast packets to hosts on a specified subnet without broadcasting those packets to hosts on the entire network.

This example shows how to enable a subnet to receive IP directed broadcast packets so you can perform backups and other network management tasks remotely:

Requirements

This example uses the following software and hardware components:

- Junos OS Release 9.4 or later for EX Series switches
- One PC
- One EX Series switch

Before you configure IP directed broadcast for a subnet:

- Ensure that the subnet does not have a direct connection to the Internet.
- Configure routed VLAN interfaces (RVIs) for the ingress and egress VLANs on the switch. See *Configuring Routed VLAN Interfaces (CLI Procedure)* or *Configuring VLANs for EX Series Switches (J-Web Procedure)*.

Overview and Topology

You might want to perform remote administration tasks such as backups and wake-on-LAN (WOL) application tasks to manage groups of clients on a subnet. One way to do this is to send IP directed broadcast packets targeted at the hosts in a particular target subnet.

The network forwards IP directed broadcast packets as if they were unicast packets. When the IP directed broadcast packet is received by a VLAN that is enabled for **targeted-broadcast**, the switch broadcasts the packet to all the hosts in its subnet.

In this topology (see Figure 7 on page 57), a host is connected to an interface on an EX Series switch to manage the clients in subnet **10.1.2.1/24**. When the switch receives a packet with the broadcast IP address of the target subnet as its destination address, it forwards the packet to the subnet's Layer 3 interface and broadcasts it to all the hosts within the subnet.

Figure 7: Topology for IP Directed Broadcast

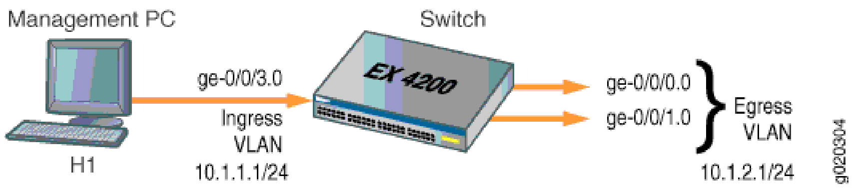

Table 11 on page 57 shows the settings of the components in this example.

Table 11: Components of the IP Directed Broadcast Topology

Property	Settings
Switch hardware	EX Series switch
Ingress VLAN name	**v0**
Ingress VLAN IP address	**10.1.1.1/24**
Egress VLAN name	**v1**
Egress VLAN IP address	**10.1.2.1/24**
Interfaces in VLAN **v0**	**ge-0/0/3.0**
Interfaces in VLAN **v1**	**ge-0/0/0.0** and **ge-0/0/1.0**

Configuration

To configure IP directed broadcast on a subnet to enable remote management of its hosts:

CLI Quick Configuration

To quickly configure the switch to accept IP directed broadcasts targeted at subnet **10.1.2.1/24**, copy the following commands and paste them into the switch's terminal window:

```
[edit]
set interfaces ge-0/0/0.0 family ethernet-switching vlan members v1
set interfaces ge-0/0/1.0 family ethernet-switching vlan members v1
set interfaces vlan.1 family inet address 10.1.2.1/24
set interfaces ge-0/0/3.0 family ethernet-switching vlan members v0
set interfaces vlan.0 family inet address 10.1.1.1/24
set vlans v1 l3-interface vlan.1
set vlans v0 l3-interface vlan.0
set interfaces vlan.1 family inet targeted-broadcast
```

Step-by-Step Procedure

To configure the switch to accept IP directed broadcasts targeted at subnet **10.1.2.1/24**:

1. Add logical interface **ge-0/0/0.0** to VLAN **v1**:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
```

2. Add logical interface **ge-0/0/1.0** to VLAN **v1**:

```
[edit interfaces]
```

```
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```

3. Configure the IP address for the egress VLAN, **v1**:

```
[edit interfaces]
user@switch# set vlan.1 family inet address 10.1.2.1/24
```

4. Add logical interface **ge-0/0/3.0** to VLAN **v0**:

```
[edit interfaces]
user@switch# set ge-0/0/3.0 family ethernet-switching vlan members v0
```

5. Configure the IP address for the ingress VLAN:

```
[edit interfaces]
user@switch# set vlan.0 family inet address 10.1.1.1/24
```

6. To route traffic between the ingress and egress VLANs, associate a Layer 3 interface with each VLAN:

```
[edit vlans]
user@switch# set v1 l3-interface (VLANs)vlan.1
user@switch# set v0 l3–interface vlan.0
```

7. Enable the Layer 3 interface for the egress VLAN to receive IP directed broadcasts:

```
[edit interfaces]
user@switch# set vlan.1 family inet targeted-broadcast
```

Results Check the results:

```
user@switch# show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members v1;
        }
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members v1;
        }
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members v0;
        }
      }
    }
  }
  vlan {
    unit 0 {
      family inet {
```



```
            targeted-broadcast;
            address 10.1.1.1/24;
          }
        }
        unit 1 {
          family inet {
            targeted-broadcast;
            address 10.1.2.1/24;
          }
        }
      }
    vlans {
      default;
      v0 {
        l3-interface vlan.0;
      }
      v1 {
        l3-interface vlan.1;
      }
    }
```

Related Documentation

- Configuring IP Directed Broadcast (CLI Procedure) on page 137

Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches

EX8200 switches support multicast load balancing on link aggregation groups (LAGs). Multicast load balancing evenly distributes Layer 3 routed multicast traffic over the LAGs, You can aggregate up to twelve 10-gigabit Ethernet links to form a 120-gigabit virtual link or LAG. The MAC client can treat this virtual link as if it were a single link to increase bandwidth, provide graceful degradation as link failures occur, and increase availability. On EX8200 switches, multicast load balancing is enabled by default. However, if it is explicitly disabled, you can reenable it. .

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

NOTE: Only EX8200 standalone switches with 10-gigabit links support multicast load balancing. Virtual Chassis does not support multicast load balancing.

This example shows how to configure a LAG and reenable multicast load balancing:

- Requirements on page 60
- Overview and Topology on page 60

- Configuration on page 61
- Verification on page 63

Requirements

This example uses the following hardware and software components:

- Two EX8200 switches, one used as the access switch and one used as the distribution switch
- Junos OS Release 12.2 or later for EX Series switches

Before you begin:

- Configure four 10-gigabit interfaces on the EX8200 distribution switch: xe-0/1/0, xe-1/1/0, xe-2/1/0, and xe-3/1/0. See "Configuring Gigabit Ethernet Interfaces (CLI Procedure)" on page 66.

Overview and Topology

Multicast load balancing uses one of seven hashing algorithms to balance traffic between the individual 10-gigabit links in the LAG. For a description of the hashing algorithms, see **multicast-loadbalance**. The default hashing algorithm is crc-sgip. You can experiment with the different hashing algorithms until you determine the one that best balances your Layer 3 routed multicast traffic.

When a link larger than 10 gigabits is needed on an EX8200 switch, you can combine up to twelve 10-gigabit links to create more bandwidth. This example uses the link aggregation feature to combine four 10-gigabit links into a 40-gigabit link on the distribution switch. In addition, multicast load balancing is enabled to ensure even distribution of Layer 3 routed multicast traffic on the 40-gigabit link. In the sample topology illustrated in Figure 8 on page 61, an EX8200 switch in the distribution layer is connected to an EX8200 switch in the access layer.

NOTE: Link speed is automatically determined based on the size of the LAG configured. For example, if a LAG is composed of four 10-gigabit links, the link speed is 40 gigabits per second).

NOTE: The default hashing algorithm, crc-sgip, involves a cyclic redundancy check of both the multicast packet source and group IP addresses.

Figure 8: 40-Gigabit LAG Composed of Four 10-Gigabit Links

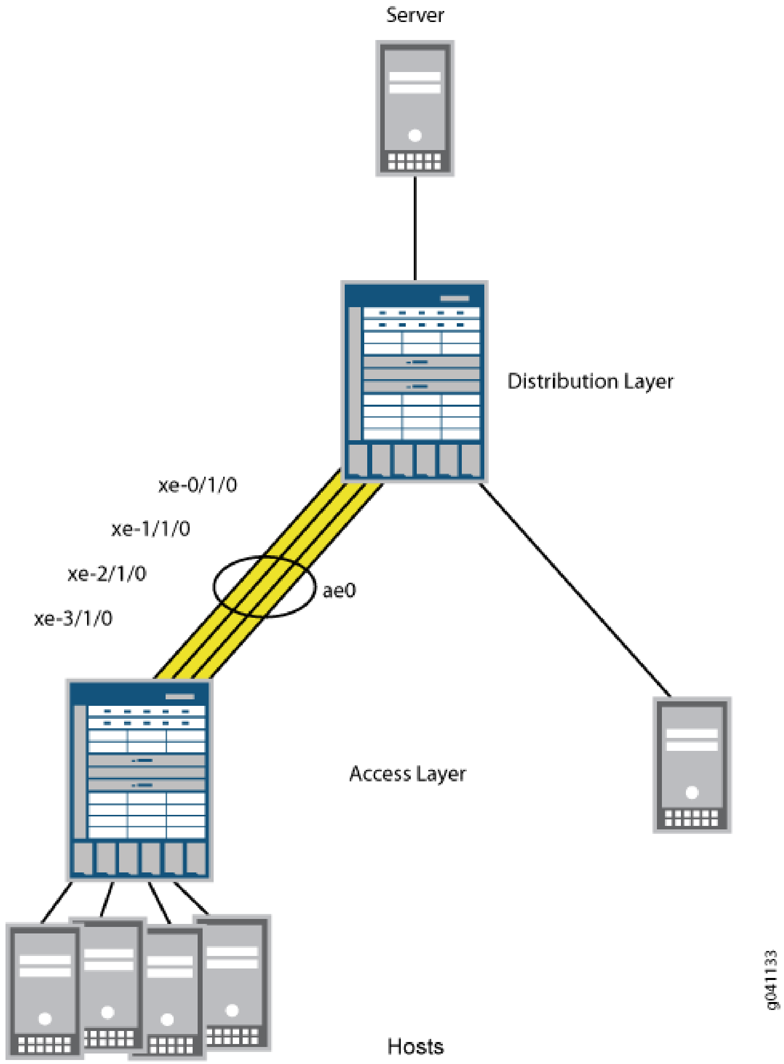

You will configure a LAG on each switch and reenable multicast load balancing. When reenabled, multicast load balancing will automatically take effect on the LAG, and the speed is set to 10 gigabits per second for each link in the LAG. Link speed for the 40-gigabit LAG is automatically set to 40 gigabits per second.

Configuration

CLI Quick Configuration

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

```
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces xe-0/1/0 ether-options 802.3ad ae0
set interfaces xe-1/1/0 ether-options 802.3ad ae0
set interfaces xe-2/1/0 ether-options 802.3ad ae0
```

```
set interfaces xe-3/1/0 ether-options 802.3ad ae0
set chassis multicast-loadbalance hash-mode crc-gip
```

Step-by-Step Procedure

To configure a LAG and reenable multicast load balancing:

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

   ```
   [edit chassis]
   user@switch# set aggregated-devices ethernet device-count 1
   ```

2. Specify the minimum number of links for the aggregated Ethernet interface (ae*x*), that is, the LAG, to be labeled **up**:

   NOTE: By default, only one link needs to be up for the LAG to be labeled **up**.

   ```
   [edit interfaces]
   user@switch# set ae0 aggregated-ether-options minimum-links 1
   ```

3. Specify the four members to be included within the LAG:

   ```
   [edit interfaces]
   user@switch# set xe-0/1/0 ether-options 802.3ad ae0
   user@switch# set xe-1/1/0 ether-options 802.3ad ae0
   user@switch# set xe-2/1/0 ether-options 802.3ad ae0
   user@switch# set xe-3/1/0 ether-options 802.3ad ae0
   ```

4. Reenable multicast load balancing:

   ```
   [edit chassis]
   user@switch# set multicast-loadbalance
   ```

   NOTE: You do not need to set link speed the way you do for LAGs that do not use multicast load balancing. Link speed is automatically set to 40 gigabits per second on a 40-gigabit LAG.

5. You can optionally change the value of the **hash-mode** option in the multicast-loadbalance statement to try different algorithms until you find the one that best distributes your Layer 3 routed multicast traffic.

   If you change the hashing algorithm when multicast load balancing is disabled, the new algorithm takes effect after you reenable multicast load balancing.

Results

Check the results of the configuration:

```
user@switch> show configuration
chassis
 aggregated-devices {
  ethernet {
     device-count 1;
  }
 }
 multicast-loadbalance {
  hash-mode crc-gip;
```



```
}

interfaces
 xe-0/1/0 {
 ether-options {
  802.3ad ae0;
 }
 }
 xe-1/1/0 {
  ether-options {
  802.3ad ae0;
 }
 }
 xe-2/1/0 {
  ether-options {
  802.3ad ae0;
 }
 }
 xe-3/1/0 {
  ether-options {
  802.3ad ae0;
 }
 }
 ae0 {
  aggregated-ether-options {
  minimum-links 1;
 }
 }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Status of a LAG Interface on page 63
- Verifying Multicast Load Balancing on page 64

Verifying the Status of a LAG Interface

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

Action Verify that the **ae0** LAG has been created:

```
user@switch> 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 interface name ae*x* indicates that this is a LAG. *A* stands for aggregated, and *E* stands for Ethernet. The number differentiates the various LAGs.

Verifying Multicast Load Balancing

Purpose Check that traffic is load-balanced equally across paths.

Action Verify load balancing across the four interfaces:

user@switch> **monitor interface traffic**

```
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
ibmoem02-re1                       Seconds: 3                  Time: 16:06:14

Interface     Link  Input packets        (pps)     Output packets        (pps)
 xe-0/1/0      Up         2058834        (10)           7345862         (19)
 xe-1/1/0      Up         2509289         (9)           6740592         (21)
 xe-2/1/0      Up         8625688        (90)          10558315         (20)
 xe-3/1/0      Up         2374154        (23)          71494375          (9)
```

Meaning The interfaces should be carrying approximately the same amount of traffic.

Related Documentation

- Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure) on page 144
- Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches on page 24

CHAPTER 3

Configuration Tasks

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

NOTE: This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

An Ethernet interface must be configured for optimal performance in a high-traffic network. EX Series switches include a factory default configuration that:

- Enables all the network interfaces on the switch
- Sets a default port mode (access)
- Sets default link settings

- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching** (except on EX8200 switches and Virtual Chassis)
- Specifies Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP)

This topic describes:

Configuring VLAN Options and Port Mode

By default, when you boot a switch and use the factory default configuration, or when you boot the switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode and accept only untagged packets from the VLAN named **default**. You can optionally configure another VLAN and use that instead of **default**. You can also configure a port to accept untagged packets from the user-configured VLAN. For details on this concept (native VLAN), see *Understanding Bridging and VLANs on EX Series Switches*

If you are connecting either a desktop phone, wireless access point or a security camera to a Power over Ethernet (PoE) port, you can configure some parameters for the PoE interface. PoE interfaces are enabled by default. For detailed information on PoE settings, see *Configuring PoE on EX Series Switches (CLI Procedure)*.

If you are connecting a device to other switches and to routers on the LAN, you need to assign the interface to a logical port and configure the logical port as a trunk port. See "Port Role Configuration with the J-Web Interface (with CLI References)" on page 77 for more information about port configuration.

If you are connecting to a server that contains virtual machines and a VEPA for packet aggregation from those virtual machines, configure the port as a tagged-access port. See *Understanding Bridging and VLANs on EX Series Switches* for more information about tagged access.

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for trunk port mode:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching port-mode trunk
```

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for tagged-access port mode:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching port-mode tagged-access
```

Configuring the Link Settings

EX Series switches include a factory default configuration that enables interfaces with the following link settings:

- All Gigabit Ethernet interfaces are set to **auto-negotiation**.
- The speed for Gigabit Ethernet interfaces is set to **auto**, allowing the interface to operate at 10m, 100m, or 1g. The link operates at the highest possible speed, depending on the capabilities of the remote end.
- The flow control for Gigabit Ethernet interfaces and 10-Gigabit Ethernet interfaces is set to **enabled**.
- The link mode is set to **auto**, allowing the interface to operate as either full duplex or half duplex. The link operates as full duplex unless this mode is not supported at the remote end.
- The 10-Gigabit Ethernet fiber interfaces default to **no auto-negotiation**. The default speed is 10g and the default link mode is full duplex.

To configure the link settings:

- Set link settings for a Gigabit Ethernet interface:

  [edit]
  user@switch# **set interfaces ge-*fpc/pic/port* ether-options**

- Set link settings for a 10-Gigabit Ethernet interface:

  [edit]
  user@switch# **set interfaces xe-*fpc/pic/port* ether-options**

NOTE: On EX Series switches, *fpc* can have the following values:

- On an EX2200 switch, an EX3200 switch, a standalone EX3300 switch, a standalone EX4200 switch, and a standalone EX4500 switch, FPC refers to the switch itself. The FPC number is always 0 on these switches.
- On an EX3300 Virtual Chassis, an EX4200 Virtual Chassis, an EX4500 Virtual Chassis, or a mixed EX4200 and EX4500 Virtual Chassis, the FPC number indicates the member ID of the switch within the Virtual Chassis.
- On an EX6200 switch and a standalone EX8200 switch, the FPC number indicates the slot number of the line card that contains the physical interface. On an EX6200 switch, the FPC number also indicates the slot number of the Switch Fabric and Routing Engine (SRE) module that contains the uplink port.
- On an EX8200 Virtual Chassis, the FPC number indicates the slot number of the line card on the Virtual Chassis. The line card slots on Virtual Chassis member 0 are numbered 0 through 15; on member 1, they are numbered 16 through 31, and so on.

pic can have the following values:

- On EX2200, EX3200, EX3300, EX4200, and EX4500 switches, the PIC number is 0 for all built-in interfaces (interfaces that are not an uplink port).
- On EX2200, EX3200, and EX4200 switches, the PIC number is 1 for uplink ports.
- On EX4500 switches, the PIC number is 1 for uplink ports on the left-hand uplink module and 2 for uplink ports on the right-hand uplink module.
- On EX6200 and EX8200 switches, the PIC number is always 0.

The **ether-options** statement allows you to modify the configuration:

- **802.3ad**—Specify an aggregated Ethernet bundle. See "Configuring Aggregated Ethernet Links (CLI Procedure)" on page 121.
- **auto-negotiation**—Enable or disable autonegotation of flow control, link mode, and speed.
- **flow-control**—Enable or disable flow control.

- link-mode—Specify **full-duplex**, **half-duplex**, or **automatic**.
- loopback—Enable or disable loopback mode.
- speed—Specify **10m**, **100m**, **1g**, or **autonegotiation**.

Configuring the IP Options

To specify an IP address for the logical unit using IPv4:

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

To specify an IP address for the logical unit using IPv6:

```
[edit]
```
`user@switch#` **set interfaces *interface-name* unit *logical-unit-number* family inet6 address *ip-address***

NOTE: Access interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by default. You might have to delete this or another user-configured family setting before changing the setting to **family inet** or **family inet6**.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- Monitoring Interface Status and Traffic on page 289
- **show interfaces ge- on page 356**
- **show interfaces xe- on page 375**
- Understanding Interface Naming Conventions on EX Series Switches on page 6

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

You can configure specific properties on your Ethernet interface to ensure optimal performance of your network in a high-traffic environment.

To configure properties on a Gigabit Ethernet interface, a 10-Gigabit Ethernet interface, and a 40-Gigabit Ethernet interface on an EX Series switch:

1. Select **Interfaces** > **Ports**.

   The page that is displayed lists Gigabit Ethernet, 10-Gigabit Ethernet interfaces, and 40-Gigabit Ethernet interfaces, and their link statuses.

   NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See *Using the Commit Options to Commit Configuration Changes (J-Web Procedure)* for details about all commit options.

2. Select the interface you want to configure. For an EX8200 Virtual Chassis configuration, select the member and the FPC slot if the interface you want to configure is not listed under **Ports** in the top table on the page.

   Details for the selected interface, such as administrative status, link status, speed, duplex, and flow control, are displayed in the **Details of port** table on the page.

   NOTE: You can select multiple interfaces and modify their settings at the same time. However, while doing this, you cannot modify the IP address or enable or disable the administrative status of the selected interfaces.

   NOTE: In the J-Web interface, you cannot configure interface ranges and interface groups.

3. Click **Edit** and select the set of options you want to configure first:

- Port Role—Enables you to assign a profile for the selected interface.

  NOTE: When you select a particular port role, preconfigured port security parameters are set for the VLAN that the interface belongs to. For example, if you select the port role **Desktop**, the port security options **examine-dhcp** and **arp-inspection** are enabled on the VLAN that the interface belongs to. If there are interfaces in the VLAN that have static IP addresses, those interfaces might lose connectivity because those static IP addresses might not be present in the DHCP pool. Therefore, when you select a port role, ensure that the corresponding port security settings for the VLAN are applicable to the interface.

  For basic information about port security features such as DHCP snooping (CLI option **examine-dhcp**) or dynamic ARP inspection (DAI) (CLI option **arp-inspection**), see *Configuring Port Security (J-Web Procedure)*. For detailed descriptions of port security features, see the Port Security topics in the EX Series documentation at http://www.juniper.net/techpubs/.

  Click **Details** to view the configuration parameters for the selected port role.

- VLAN—Enables you to configure VLAN options for the selected interface.
- Link—Enables you to modify the following link options for the selected interface:
  - Speed
  - MTU
  - Autonegotiation
  - Flow Control
  - Duplex
  - Media Type
- IP—Enables you to configure an IP address for the interface.

4. Configure the interface by configuring options in the selected option set. See Table 12 on page 73 for details of the options.
5. Repeat Steps 3 and 4 for the remaining option sets that you want to configure for the interface.

   NOTE: To enable or disable the administrative status of a selected interface, click **Enable Port** or **Disable Port**.

Table 12: Port Edit Options

Field	Function	Your Action
Port Role Options		
Port Role	Specifies a profile (role) to assign to the interface. NOTE: After a port role is configured on the interface, you cannot specify VLAN options or IP options. NOTE: Port roles are not supported by the et interfaces (40-Gigabit Ethernet interfaces) on EX4300 and EX4550 switches. NOTE: Only the following port roles can be applied on EX8200 switch interfaces: • Default • Layer 2 uplink • Routed uplink	
Default	Applies the default role. The interface family is set to **ethernet-switching**, port mode is set to **access**, and RSTP is enabled.	1. Click **Details** to view CLI commands for this role. 2. Click **OK**.
Desktop	Applies the desktop role. The interface family is set to **ethernet-switching**, port mode is set to **access**, RSTP is enabled with the **edge** and **point-to-point** options, and port security parameters (MAC limit =1; dynamic ARP inspection and DHCP snooping enabled) are set.	1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.
Desktop and Phone	Applies the desktop and phone role. The interface family is set to **ethernet-switching**, port mode is set to **access**, port security parameters (MAC limit =1; dynamic ARP Inspection and DHCP snooping enabled) are set, and recommended class-of-service (CoS) parameters are specified for forwarding classes, schedulers, and classifiers. See Table 13 on page 76 for more CoS information.	1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. You can also select an existing VoIP VLAN configuration or a new VoIP VLAN configuration to be associated with the interface. NOTE: VoIP is not supported on EX8200 switches. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.
Wireless Access Point	Applies the wireless access point role. The interface family is set to **ethernet-switching**, port mode is set to **access**, and RSTP is enabled with the **edge** and **point-to-point** options.	1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. Type the VLAN ID for a new VLAN. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.

Table 12: Port Edit Options *(continued)*

Field	Function	Your Action
Routed Uplink	Applies the routed uplink role. The interface family is set to **inet**, and recommended CoS parameters are set for schedulers and classifiers. See Table 13 on page 76 for more CoS information.	To specify an IPv4 address: 1. Select the **IPv4 address** check box. 2. Type an IP address—for example: **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. 4. Click **OK**. To specify an IPv6 address: 1. Select the **IPv6 address** check box. 2. Type an IP address—for example: **2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix. 4. Click **OK**. NOTE: IPv6 is not supported on EX2200 VC switches.
Layer 2 Uplink	Applies the Layer 2 uplink role. The interface family is set to **ethernet-switching**, port mode is set to **trunk**, RSTP is enabled with the **point-to-point** option, and trusted DHCP is configured for port security.	1. For this port role, you can select a VLAN member and associate a native VLAN with the interface. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.
None	Specifies that no port role is configured for the selected interface.	
NOTE: For an EX8200 switch, dynamic ARP inspection and DHCP snooping parameters are not configured.		
VLAN Options		

Table 12: Port Edit Options *(continued)*

Field	Function	Your Action
Port Mode	Specifies the mode of operation for the interface: trunk or access.	If you select **Trunk**, you can: 1. Click **Add** to add a VLAN member. 2. Select the VLAN and click **OK**. 3. (Optional) Associate a native VLAN with the interface. 4. Click **OK**. If you select **Access**, you can: 1. Select the VLAN member to be associated with the interface. 2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN. NOTE: VoIP is not supported on EX8200 switches. 3. Click **OK**.
Link Options		
MTU (bytes)	Specifies the maximum transmission unit size (MTU) for the interface.	Type a value from **256** through **9216**. The default MTU size for Gigabit Ethernet interfaces is **1514**.
Speed	Specifies the speed for the mode.	Select one of the following values: **10 Mbps**, **100 Mbps**, **1000 Mbps**, or **Auto-Negotiation.** NOTE: EX4300 switches supports **Auto-Negotiation 10M-100M** apart from the values mentioned above.
Duplex	Specifies the link mode.	Select one: **automatic**, **half**, or **full**. NOTE: Link mode **half** is not supported on EX4300 switches.
Description	Describes the link. NOTE: If the interface is part of a link aggregation group (LAG), only the **Description** option is enabled. Other Port Edit options are unavailable.	Enter a brief description for the link.
Enable Auto Negotiation	Enables or disables autonegotiation.	Select the check box to enable autonegotiation, or clear the check box to disable it. By default, autonegotiation is enabled.
Enable Flow Control	Enables or disables flow control.	Select the check box to enable flow control to regulate the amount of traffic sent out of the interface, or clear the check box to disable flow control and permit unrestricted traffic. Flow control is enabled by default.

Table 12: Port Edit Options *(continued)*

Field	Function	Your Action
Media Type	Specifies the media type selected.	Select the check box to enable the media type. Then select **Copper** or **Fiber**.
IP Options		
IPv4 Address	Specifies an IPv4 address for the interface. NOTE: If the IPv4 Address check box is cleared, the interface still belongs to the **inet** family.	1. Select the **IPv4 address** check box to specify an IPv4 address. 2. Type an IP address—for example: **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. 4. Click **OK**.
IPv6 Address	Specifies an IPv6 address for the interface. NOTE: If the IPv6 Address check box is cleared, the interface still belongs to the **inet** family.	1. Select the **IPv6 address** check box to specify an IPv6 address. 2. Type an IP address—for example: **2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix. 4. Click **OK**. NOTE: IPv6 address is not supported on EX2200 and EX4500 switches.

Table 13: Recommended CoS Settings for Port Roles

CoS Parameter	Recommended Settings
Forwarding Classes	There are four forwarding classes: • **voice**—Queue number is set to 7. • **expedited-forwarding**—Queue number is set to 5. • **assured-forwarding**—Queue number is set to 1. • **best-effort**—Queue number is set to 0.
Schedulers	The schedulers and their settings are: • Strict-priority—Transmission rate is set to 10 percent and buffer size to 5 percent. • Expedited-scheduler—Transmission rate is set to 30 percent, buffer size to 30 percent, and priority to **low**. • Assured-scheduler—Transmission rate is set to 25 percent, buffer size to 25 percent, and priority to **low**. • Best-effort scheduler—Transmission rate is set to 35 percent, buffer size to 40 percent, and priority to **low**.
Scheduler maps	When a desktop and phone, routed uplink, or Layer 2 uplink role is applied on an interface, the forwarding classes and schedulers are mapped using the scheduler map.
ieee-802.1 classifier	Imports the default **ieee-802.1** classifier configuration and sets the loss priority to **low** for the code point 101 for the **voice** forwarding class.

Table 13: Recommended CoS Settings for Port Roles *(continued)*

CoS Parameter	Recommended Settings
dscp classifier	Imports the default **dscp** classifier configuration and sets the loss priority to **low** for the code point 101110 for the **voice** forwarding class.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Monitoring Interface Status and Traffic on page 289
- EX Series Switches Interfaces Overview on page 3
- *Junos OS CoS for EX Series Switches Overview*
- Understanding Interface Naming Conventions on EX Series Switches on page 6

Port Role Configuration with the J-Web Interface (with CLI References)

When you configure Gigabit Ethernet interface properties with the J-Web interface (Configure > Interfaces) you can optionally select pre-configured port roles for those interfaces. When you select a role from the **Port Role** field and apply it to a port, the J-Web interface modifies the switch configuration using CLI commands. Table 14 on page 77 lists the CLI commands applied for each port role.

NOTE: If there is an existing port role configuration, it is cleared before the new port role configuration is applied.

Table 14: Port Role Configuration Summary

Configuration Description	CLI Commands
Default Port Role	
Set the port role to **Default**.	`set interfaces` *`interface`*`apply-macro juniper-port-profile Default`
Set port family to **ethernet-switching**. Set port mode to **access**.	`set interfaces` *`interface`* `unit 0 family ethernet-switching port-mode access`
Enable RSTP if redundant trunk groups are not configured.	`delete protocols rstp interface` *`interface`* `disable`
Disable RSTP if redundant trunk groups are configured.	`set protocols rstp interface` *`interface`* `disable`
Desktop Port Role	

Table 14: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands
Set the port role to desktop.	`set interfaces` *interface* `apply-macro juniper-port-profile Desktop`
Set VLAN if new VLAN is specified.	set vlans <*vlan name*> vlan-id <*vlan-id*>
Set port family to **ethernet-switching**. Set Port Mode to **Access**.	`set interfaces` *interface* `unit 0 family ethernet-switching port-mode access`
Set VLAN if new VLAN is specified.	`set interfaces` *interface* `unit 0 family ethernet-switching vlan members` *vlan-members*
Set port security parameters.	`set ethernet-switching-options secure-access-port vlan MacTest arp-inspection`
Set RSTP protocol with **edge** option.	`set protocols rstp interface` *interface* `edge`
RSTP protocol is disabled if redundant trunk groups are configured.	`set protocols rstp interface` *interface* `disable`
Desktop and Phone Port Role	
Set the port role to desktop and phone.	`set interfaces` *interface*`apply-macro juniper-port-profile Desktop and Phone`
Set data VLAN if new VLAN is specified. Set voice VLAN if new voice VLAN is specified.	`set vlans` *vlan-name*`vlan-id` *vlan id*
Set port family to **ethernet-switching.** Set Port Mode to **access**.	`set interfaces` *interface*`unit 0 family ethernet-switching port-mode access`
Set data VLAN on port stanza.	`set interfaces` *interface* `unit 0 family ethernet-switching vlan members` *vlan-members*
Set port security parameters.	`set ethernet-switching-options secure-access-port vlan MacTest arp-inspection`
Set VOIP VLAN.	`set ethernet-switching-options voip interface interface.0 vlan vlan` *vlan name*

Set class of service parameters
 	
SCHEDULER_MAP=juniper-port-profile-map
 	
IEEE_CLASSIFIER=juniper-ieee-classifier
 	
DSCP_CLASSIFIER=juniper-dscp-classifier
	`set class-of-service interfaces` *interface*`scheduler-map juniper-port-profile-map`
	`set class-of-service interfaces` *interface* `unit 0 classifiers ieee-802.1 juniper_ieee_classifier`
	`set class-of-service interfaces` *interface*`unit 0 classifiers dscp juniper-dscp-classifier`

Set CoS Configuration	Refer Table 15 on page 80 for details.
Wireless Access Point Port Role	
Set the port role to wireless access point.	`set interfaces` *interface* `apply-macro juniper-port-profile Wireless Access Point`

Table 14: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands
Set VLAN on VLANs stanza.	`set vlans` *`vlan name`*`vlan-id` *`vlan-id`*
Set port family to **ethernet-ewitching** Set port mode to **Access**.	`set interfaces` *`interface`* `unit 0 family ethernet-switching port-mode access`
Set VLAN on port stanza.	`set interfaces` *`interface`* `unit 0 family ethernet-switching vlan members` *`vlan-members`*
Set RSTP protocol with edge option.	`set protocols rstp interface` *`interface`* `edge`
RSTP protocol is disabled if redundant trunk groups are configured.	`set protocols rstp interface` *`interface`* `disable`
Routed Uplink Port Role	
Set the port role to Routed Uplink.	`set interfaces` *`interface`* `apply-macro juniper-port-profile Routed Uplink`
Set port family to inet. Set IP address on the port.	`set interfaces` *`interface`*`unit 0 family inet address` *`ipaddress`*

Set class-of-service parameters
 	
SCHEDULER_MAP=juniper-port-profile-map
 	
IEEE_CLASSIFIER=juniper-ieee-classifier
 	
DSCP_CLASSIFIER=juniper-dscp-classifier
	`set class-of-service interfaces` *`interface`*`scheduler-map juniper-port-profile-map`
	`set class-of-service interfaces` *`interface`* `unit 0 classifiers ieee-802.1 juniper_ieee_classifier`
	`set class-of-service interfaces` *`interface`*`unit 0 classifiers dscp juniper-dscp-classifier`

Set CoS configuration	Refer Table 15 on page 80 for details.
Layer 2 Uplink Port Role	
Set the port role to **Layer 2 Uplink**.	`set interfaces` *`interface`* `apply-macro juniper-port-profile Layer2 Uplink`
Set port family to **ethernet-switching** Set port mode to **trunk**.	`set interfaces` *`interface`* `unit 0 family ethernet-switching port-mode trunk`
Set Native VLAN name.	`set interfaces` *`interface`* `unit 0 family ethernet-switching native-vlan-id` *`vlan-name`*
Set the port as part of all valid VLANs; "valid" refers to all VLANs except native VLAN and voice VLANs.	`set interfaces` *`interface`* `unit 0 family ethernet-switching vlan members` *`vlan-members`*
Set port security parameter.	`set ethernet-switching-options secure-access-port dhcp-trusted`
Set RSTP protocol with point-to-point option.	`set protocols rstp interface` *`interface`* `mode point-to-point`
Disable RSTP if redundant trunk groups are configured.	`set protocols rstp interface` *`interface`* `disable`

Table 14: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands

Set class-of-service parameters.
 	
SCHEDULER_MAP=juniper-port-profile-map
 	
IEEE_CLASSIFIER=juniper_ieee_classifier
 	
DSCP_CLASSIFIER=juniper_dscp_classifier
	`set class-of-service interfaces` *`interface`*`scheduler-map juniper-port-profile-map`
	`set class-of-service interfaces` *`interface`* `unit 0 classifiers ieee-802.1 juniper_ieee_classifier`
	`set class-of-service interfaces` *`interface`*`unit 0 classifiers dscp juniper-dscp-classifier`

Set CoS configuration	Refer to Table 15 on page 80 for details.

Table 15 on page 80 lists the CLI commands for the recommended CoS settings that are committed when the CoS configuration is set.

Table 15: Recommended CoS Settings for Port Roles

CoS Parameter	CLI Command
Forwarding Classes	
voice	`set class-of-service forwarding-classes class voice queue-num 7`
expedited-forwarding	`set class-of-service forwarding-classes class expedited-forwarding queue-num 5`
assured-forwarding	`set class-of-service forwarding-classes class assured-forwarding queue-num 1`
best-effort	`set class-of-service forwarding-classes class best-effort queue-num 0`
Schedulers	

strict-priority-scheduler	The CLI commands are:
	
	• `set class-of-service schedulers strict-priority-scheduler transmit-rate percent 10`
	`set class-of-service schedulers strict-priority-scheduler buffer-size percent 5`
	`set class-of-service schedulers strict-priority-scheduler priority strict-high`

expedited-scheduler	The CLI commands are:
	
	• `set class-of-service schedulers expedited-scheduler transmit-rate percent 30`
	`set class-of-service schedulers expedited-scheduler buffer-size percent 30`
	`set class-of-service schedulers expedited-scheduler priority low`

assured-scheduler	The CLI commands are:
	
	`set class-of-service schedulers assured-scheduler transmit-rate percent 25`
	`set class-of-service schedulers strict-priority-scheduler buffer-size percent 25`
	`set class-of-service schedulers strict-priority-scheduler priority low`

Table 15: Recommended CoS Settings for Port Roles *(continued)*

CoS Parameter	CLI Command

best-effort-scheduler	The CLI commands are:
	
	`set class-of-service schedulers best-effort-scheduler transmit-rate percent 35`
	`set class-of-service schedulers best-effort-scheduler buffer-size percent 40`
	`set class-of-service schedulers best-effort-scheduler priority low`

Classifiers	The classifiers are:
	
	`set class-of-service classifiers ieee-802.1 juniper_ieee_classifier import default forwarding-class voice loss-priority low code-points 101`
	`set class-of-service classifiers dscp juniper_dscp_classifier import default forwarding-class voice loss-priority low code-points 101110`

Related Documentation

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

Adding a Logical Unit Description to the Configuration

You can include a text description of each logical unit in the configuration file. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no impact on the interface's configuration. To add a text description, include the **description** statement:

description *text*;

You can include this statement at the following hierarchy levels:

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

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.

NOTE: You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See "*Using DHCP Relay Agent Option 82 Information*" in the *Junos OS Subscriber Management and Services Library*.

For information about describing physical interfaces, see *Configuring Interface Description*.

Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
disable;
```

CAUTION: Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.

NOTE: On the router, when you use the **disable** statement at the **edit interfaces** hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet PICs with SFP and XFP transceivers do support it and the laser will be turned off when the interface is disabled.

Table 16: Effect of set interfaces disable <interface_name> on T series PICs

PIC Model Number	PIC Description	Type of PIC	Behaviour
PF-12XGE-SFPP	10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 Router)	5	Tx laser disabled
PF-24XGE-SFPP	10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (T4000 Router)	5	Tx laser disabled
PF-1CGE-CFP	100-Gigabit Ethernet PIC with CFP (T4000 Router)	5	Tx laser disabled
PD-4XGE-XFP	10-Gigabit Ethernet, 4-port LAN/WAN XFP	4	Tx laser disabled
PD-5-10XGE-SFPP	10-Gigabit LAN/WAN with SFP+	4	Tx laser disabled
PD-1XLE-CFP	40-Gigabit with CFP	4	Tx laser disabled
PD-1CE-CFP-FPC4	100-Gigabit with CFP	4	Tx laser disabled
PD-TUNNEL	40-Gigabit Tunnel Services	4	NA
PD-4OC192-SON-XFP	OC192/STM64, 4-port XFP	4	Tx laser not disabled
PD-1OC768-SON-SR	OC768c/STM256, 1-port	4	Tx laser not disabled

WARNING: Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

Example: Disabling a Physical Interface

Sample interface configuration:

```
[edit interfaces]
user@host# show
ge-0/3/2 {
    unit 0 {
        description CE2-to-PE1;
        family inet {
            address 20.1.1.6/24;
        }
    }
}
```

Disabling the interface:

```
[edit interfaces]
user@host# set ge-0/3/2 disable
```

Verifying the interface configuration:

```
[edit interfaces]
user@host# show
ge-0/3/2 {
    disable; # Interface is marked as disabled.
    unit 0 {
        description CE2-to-PE1;
        family inet {
            address 20.1.1.6/24;
        }
    }
}
```

Disabling a Logical Interface

You can unconfigure a logical interface, effectively disabling that interface, without removing the logical interface configuration statements from the configuration. To do this, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

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

When an interface is disabled, a route (pointing to the reserved target “**REJECT**”) with the IP address of the interface and a 32–bit subnet mask is installed in the routing table. See *Routing Protocols*.

Configuring Flow Control

By default, the router or switch imposes flow control to regulate the amount of traffic sent out on a Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interface. Flow control is not supported on the 4-port Fast Ethernet PIC. This is useful if the remote side of the connection is a Fast Ethernet or Gigabit Ethernet switch.

You can disable flow control if you want the router or switch to permit unrestricted traffic. To disable flow control, include the **no-flow-control** statement:

```
no-flow-control;
```

To explicitly reinstate flow control, include the **flow-control** statement:

```
flow-control;
```

You can include these statements at the following hierarchy levels:

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

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Related Documentation

- flow-control on page 213
- *Ethernet Interfaces Overview*
- EX Series Switches Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 86
- Configuring Interface IPv6 Addresses on page 89

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 atleast one of them will have 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 fec0:1:1:1::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 the Interface Bandwidth

By default, the Junos OS uses the physical interface's speed for the MIB-II object, **ifSpeed**. You can configure the logical unit to populate the **ifSpeed** variable by configuring a bandwidth value for the logical interface. The **bandwidth** statement sets an informational-only parameter; you cannot adjust the actual bandwidth of an interface with this statement.

NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

cost = reference-bandwidth/bandwidth,

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

To configure the bandwidth value for a logical interface, include the **bandwidth** statement:

```
bandwidth rate;
```

You can include this statement at the following hierarchy levels:

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

rate is the peak rate, in bps or cps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation **c**; values expressed in cells per second are converted to bits per second using the formula 1 cps = 384 bps. The value can be any positive integer. The **bandwidth** statement is valid for all logical interfaces, except multilink interfaces.

Configuring the Media MTU

The media maximum transmission unit (MTU) is the largest data unit that can be forwarded without fragmentation.

This topic contains the following sections:

Media MTU Overview

The default media MTU size used on a physical interface depends on the encapsulation used on that interface. In some cases, the default IP Protocol MTU depends on whether the protocol used is IP version 4 (IPv4) or International Organization for Standardization (ISO).

The default media MTU is calculated as follows:

Default media MTU = Default IP MTU + encapsulation overhead

When you are configuring point-to-point connections, the MTU sizes on both sides of the connections must be the same. Also, when you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size. For details about encapsulation overhead, see "Encapsulation Overhead by Encapsulation Type" on page 93.

NOTE: The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the media MTU. For example, the media MTU for a Gigabit Ethernet Version 2 interface is specified as 1514 bytes, but the largest possible frame size is actually 1518 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

The physical MTU for Ethernet interfaces does not include the 4-byte frame check sequence (FCS) field of the Ethernet frame.

A SONET/SDH interface operating in concatenated mode has a "c" added to the rate descriptor. For example, a concatenated OC48 interface is referred to as OC48c.

If you do not configure an MPLS MTU, the Junos OS derives the MPLS MTU from the physical interface MTU. From this value, the software subtracts the encapsulation-specific overhead and space for the maximum number of labels that might be pushed in the Packet Forwarding Engine. Currently, the software provides for three labels of four bytes each, for a total of 12 bytes.

In other words, the formula used to determine the MPLS MTU is the following:

MPLS MTU = physical interface MTU – encapsulation overhead – 12

If you configure an MTU value by including the **mtu** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family mpls]** hierarchy level, the configured value is used.

How to Configure the Media MTU

To modify the default media MTU size for a physical interface, include the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level:

[edit interfaces *interface-name*]
mtu *bytes*;

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead.

NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

You configure the protocol MTU by including the **mtu** statement 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*]**

Because tunnel services interfaces are considered logical interfaces, you cannot configure the MTU setting for the physical interface. This means you cannot include the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level for the following interface types: generic routing encapsulation (**gr-**), IP-IP (**ip-**), loopback (**lo-**), link services (**ls-**), multilink services (**ml-**), and multicast (**pe-**, **pd-**). You can, however, configure the protocol MTU on tunnel interfaces, as described in “Setting the Protocol MTU” on page 103.

Encapsulation Overhead by Encapsulation Type

Table 17: Encapsulation Overhead by Encapsulation Type

Interface Encapsulation	Encapsulation Overhead (Bytes)
802.1Q/Ethernet 802.3	21
802.1Q/Ethernet Subnetwork Access Protocol (SNAP)	26
802.1Q/Ethernet version 2	18
ATM Cell Relay	4
ATM permanent virtual connection (PVC)	12
Cisco HDLC	4
Ethernet 802.3	17
Ethernet circuit cross-connect (CCC) and virtual private LAN service (VPLS)	4
Ethernet over ATM	32
Ethernet SNAP	22
Ethernet translational cross-connect (TCC)	18
Ethernet version 2	14
Extended virtual local area network (VLAN) CCC and VPLS	4
Extended VLAN TCC	22
Frame Relay	4
PPP	4
VLAN CCC	4
VLAN VPLS	4
VLAN TCC	22

Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers

Table 18: Media MTU Sizes by Interface Type for M5 and M7i Routers with CFEB, M10 and M10i Routers with CFEB, and M20 and M40 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	9192	1500
E3/T3	4474	9192	4470
Fast Ethernet	1514	1533 (4-port) 1532 (8-port) 1532 (12-port) NOTE: The maximum MTU for two 100Base-TX Fast Ethernet port FIC is 9192 bytes.	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192 NOTE: The maximum MTU for one Gigabit Ethernet port FIC is 9192 bytes.	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470

Media MTU Sizes by Interface Type for M40e Routers

Table 19: Media MTU Sizes by Interface Type for M40e Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470

Table 19: Media MTU Sizes by Interface Type for M40e Routers *(continued)*

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
E1/T1	1504	4500	1500
E3/T3	4474	4500 9192 (4-port)	4470
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	1533	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port) 9192 (4-port)	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (4-port OC3) 9192 (4-port OC3c) 4500 (1-port OC12) 4500 (4-port OC12) 4500 (4-port OC12c) 4500 (1-port OC48) 9192 (2-port OC3) 9192 (2-port OC3c) 9192 (1-port OC12c) 9192 (1-port OC48c) 4500 (1-port OC192) 9192 (1-port OC192c)	4470

Media MTU Sizes by Interface Type for M160 Routers

Table 20: Media MTU Sizes by Interface Type for M160 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500	4470
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	1533	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port) 4500 (4-port)	1500 (IPv4), 1497 (ISO)
Serial	1504	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (1- or 2-port) 4500 (4-port)	4470

Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers

Table 21: Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
Channelized DS3 IQ	4471	4500	4470
Channelized E1 IQ	1504	4500	1500
Channelized OC12 IQ	4474	9192	4470

Table 21: Media MTU Sizes by Interface Type for M7i Routers with CFEB-E, M10i Routers with CFEB-E, and M320 and M120 Routers *(continued)*

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
E1	1504	4500	1500
E3 IQ	4471	4500	4470
Fast Ethernet	1514	1533 (4-port) 1532 (8-, 12- and 48-port)	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
T1	1504	4500	1500
CT3 IQ (excluding M120)	4474	9192	4470

Media MTU Sizes by Interface Type for MX Series Routers

Table 22: Media MTU Sizes by Interface Type for MX Series Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
10-Gigabit Ethernet	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
Multi-Rate Ethernet	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
Tri-Rate Ethernet	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)

Table 22: Media MTU Sizes by Interface Type for MX Series Routers *(continued)*

Channelized SONET/SDH OC3/STM1 (Multi-Rate)	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
DS3/E3 (Multi-Rate)	1514	9192	1500 (IPv4), 1488 (MPLS), 1497 (ISO)

Media MTU Sizes by Interface Type for T320 Routers

Table 23: Media MTU Sizes by Interface Type for T320 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM	4482	9192	4470
ATM2 IQ	4482	9192	4470
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
Fast Ethernet	1514	1533 (4-port) 1532 (12- and 48-port)	1500 (IPv4), 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
CT3 IQ	4474	9192	4470

Media MTU Sizes by Interface Type for T640 Platforms

Table 24: Media MTU Sizes by Interface Type for T640 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
48-port Fast Ethernet	1514	1532	1500 (IPv4), 1497 (ISO)

Table 24: Media MTU Sizes by Interface Type for T640 Platforms *(continued)*

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
SONET/SDH	4474	9192	4470
CT3 IQ	4474	9192	4470

Media MTU Sizes by Interface Type for J2300 Platforms

Table 25: Media MTU Sizes by Interface Type for J2300 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Fast Ethernet (10/100)	1514	9192	1500
G.SHDSL	4482	9150	4470
ISDN BRI	1504	4092	1500
Serial	1504	9150	1500
T1 or E1	1504	9150	1500

Media MTU Sizes by Interface Type for J4300 and J6300 Platforms

Table 26: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ADSL2+ PIM	4482	9150	4470
Dual-port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-port Serial PIM	1504	9150	1500
Dual-port T1 or E1 PIM	1504	9150	1500
Dual-port Channelized T1/E1 PIM (channelized to DS0s)	1504	4500	1500

Table 26: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms *(continued)*

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Dual-port Channelized T1/E1 PIM (clear channel T1 or E1)	1504	9150	1500
Fast Ethernet (10/100) built-in interface	1514	9192	1500
G.SHDSL PIM	4482	9150	4470
4-port ISDN BRI PIM	1504	4092	1500
T3 (DS3) or E3 PIM	4474	9192	4470

Media MTU Sizes by Interface Type for J4350 and J6350 Platforms

Table 27: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
4-port ISDN BRI PIM	1504	4092	1500
ADSL2+ PIM	4482	9150	4470
Dual-port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-port Serial PIM	1504	9150	1500
Dual-port T1 or E1 PIM	1504	9150	1500
Dual-port Channelized T1/E1 PIM (channelized to DS0s)	1504	4500	1500
Dual-port Channelized T1/E1 PIM (clear channel T1 or E1)	1504	9150	1500
4-port Fast Ethernet (10/100) ePIM	1518	1518	1500

Table 27: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms *(continued)*

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet (10/100/1000) built-in interface	1514	9018	1500
Gigabit Ethernet (10/100/1000) Enhanced Physical Interface Module (ePIM)	1514	9018	1500
Gigabit Ethernet (10/100/1000) SFP ePIM	1514	9018	1500
G.SHDSL PIM	4482	9150	4470
T3 (DS3) or E3 PIM	4474	9192	4470

NOTE: On Gigabit Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 9018 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 9018 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: ge-0/0/0: Illegal media change. MTU invalid: 9192. Max MTU
supported on this PIC: 9018
```

On 4-port Fast Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 1518 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 1518 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: fe-3/0/1: Illegal media change. MTU invalid: 9192. Max MTU
supported on this PIC: 1518
```


Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers

Table 28: Media MTU Sizes by Interface Type for EX Series Switches and ACX Series Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)
10-Gigabit Ethernet	1514	9192	1500 (IPv4), 1497 (ISO)

NOTE: On ACX Series routers, you can configure the protocol MTU by including the **mtu** statement at the **[edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet]** or **[edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet6]** hierarchy level.

- If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
- If you are configuring the protocol MTU for both **inet** and **inet6** families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for **inet** and **inet6** families that are configured on the same logical interface.

Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers

Table 29: Media MTU Sizes by Interface Type for PTX Series Packet Transport Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
10-Gigabit Ethernet	1514	9500	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
40-Gigabit Ethernet	1514	9500	1500 (IPv4), 1488 (MPLS), 1497 (ISO)
100-Gigabit Ethernet	1514	9500	1500 (IPv4), 1488 (MPLS), 1497 (ISO)

Related Documentation

- *Configuring Interface Encapsulation on Physical Interfaces*

- Setting the Protocol MTU on page 103

Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see "Configuring the Media MTU" on page 90.

To modify the MTU for a particular protocol family, include the **mtu** statement:

```
mtu bytes;
```

You can include this statement 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*]**

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see Table 17 on page 93. If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce the protocol MTU size. (You configure the media MTU by including the **mtu** statement at the **[edit interfaces *interface-name*]** hierarchy level, as discussed in "Configuring the Media MTU" on page 90.)

NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

Related Documentation

- Configuring the Media MTU on page 90

Interface Ranges

NOTE: This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Interface Ranges*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

The Junos OS allows you to group a range of identical interfaces into an *interface range*. You first specify the group of identical interfaces in the interface range. Then you can apply a common configuration to the specified interface range, reducing the number of configuration statements required and saving time while producing a compact configuration.

Configuring Interface Ranges

To configure an interface range, include the **interface-range** statement at the **[edit interfaces]** hierarchy level.

The **interface-range** statement accepts only physical networking interface names in its definition. The following interface types are supported and example CLI descriptors are shown:

- ATM—**at-*fpc/pic/port***
- Channelized—**(coc | cstm)*n-fpc/pic/port***
- DPC—**xe-*fpc/pic/port***
- E1/E3—**(e1 | e3)-*fpc/pic/port***
- Ethernet—**(xe | ge | fe)-*fpc/pic/port***
- ISDN—**isdn-*fpc/pic/port***
- Serial—**se-*fpc/pic/port***
- SONET/SDH—**so-*fpc/pic/port***
- T1/T3—**(t1 | t3)-*fpc/pic/port***

Interfaces can be grouped either as a range of interfaces or using a number range under the **interface-range** statement definition.

Interfaces in an **interface-range** definition can be added as part of a member range or as individual members or multiple members using a number range.

To specify a member range, use the **member-range** statement at the **[edit interfaces interface-range *name*]** hierarchy level.

To specify interfaces in lexical order, use the **member-range *start-range* to *end-range*** statement.

A range for a member statement should contain the following:

- *****—All, specifies sequential interfaces from 0 through 47.

> CAUTION: The wildcard * in a member statement does not take into account the interface numbers supported by a specific interface type. Irrespective of the interface type, * includes interface numbers ranging from 0 through 47 to the interface group. Therefore, use * in a member statement with caution.

- **num**—Number, specifies one specific interface by its number.
- **[low-high]**—Numbers between low to high, specifies a range of sequential interfaces.
- **[num1, num2, num3]**—Numbers **num1**, **num2**, and **num3** specify multiple specific interfaces.

Example: Specifying an Interface Range Member Range

```
member-range ge-0/0/0 to ge-4/0/40;
```

To specify one or multiple members, use the **member** statement at the **[edit interfaces interface-range *name*]** hierarchy level.

To specify the list of interface range members individually or for multiple interfaces using regex, use the **member *list of interface names*** statement.

Example: Specifying an Interface Range Member

```
member ge-0/0/0;
member ge-0/*/*
member ge-0/[1-10]/0;
member ge-0/[1,2,3]/3;
```

Regex or wildcards are not supported for interface-type prefixes. For example, prefixes **ge**, **fe**, and **xe** must be mentioned explicitly.

An **interface-range** definition can contain both **member** and **member-range** statements within it. There is no maximum limit on the number of **member** or **member-range** statements within an interface-range. However, at least one **member** or **member-range** statement must exist within an **interface-range** definition.

Example: Interface Range Common Configuration

Configuration common to an interface range can be added as a part of the **interface-range** definition, as follows:

```
[edit]
interfaces {
```



```
    + interface-range foo {
    + member-range ge-1/0/0 to ge-4/0/40;
    + member ge-0/1/1;
    + member ge-5/[1-10]/*;
      /*Common configuration is added as part of interface-range definition*/
      mtu 256;
      hold-time up 10;
      ether-options {
        flow-control;
        speed {
          100m;
        }
        802.3ad primary;
      }
    }
  }
```

An **interface-range** definition having just **member** or **member-range** statements and no common configurations statements is valid.

These defined interface ranges can be used in other configuration hierarchies, in places where an **interface** node exists.

Example: Interface-Range foo Used Under the Protocols Hierarchy

```
protocols {
  dot1x {
    authenticator {
      interface foo{
        retries 1;
      }
    }
  }
}
```

foo should be an **interface-range** defined at the **[interfaces]** hierarchy level. In the above example, the **interface** node can accept both individual interfaces and interface ranges.

TIP: To view an interface range in expanded configuration, use the **(show | display inheritance)** command. For more information, see the *CLI User Guide*.

By default, **interface-range** is not available to configure in the CLI where the **interface** statement is available. The following locations are supported; however, some of the hierarchies shown in this list are product specific:

- **protocols dot1x authentication interface**
- **protocols dvmrp interface**
- **protocols oam ethernet lmi interface**
- **protocols esis interface**
- **protocols igmp interface**
- **protocols igmp-host client *num* interface**

- protocols mld-host client *num* interface
- protocols router-advertisement interface
- protocols isis interface
- protocols ldp interface
- protocols oam ethernet link-fault-management interface
- protocols lldp interface
- protocols link-management peer lmp-control-channel interface
- protocols link-management peer control-channel
- protocols link-management te-link *name* interface
- protocols mld interface
- protocols ospf area *id* interface
- protocols pim interface
- protocols router-discovery interface
- protocols rip group *name* neighbour
- protocols ripng group *name* neighbour
- protocols rsvp interface
- protocols snmp interface
- protocols layer2-control bpdu-block interface
- protocols layer2-control mac-rewrite interface
- protocols mpls interface
- protocols stp interface
- protocols rstp interface
- protocols mstp interface
- protocols vstp interface
- protocols mstp msti *id* interface
- protocols mstp msti vlan *id* interface
- protocols vstp vlan *name* interface
- protocols gvrp interface
- protocols igmp-snooping vlan *name* interface
- protocols lldp interface
- protocols lldp-med interface
- protocols sflow interfaces
- ethernet-switching-options analyzer *name* input [egress | ingress] interface

- **ethernet-switching-options analyzer *name* output interface**
- **ethernet-switching-options secure-access-port interface**
- **ethernet-switching-options interfaces ethernet-switching-options voip interface**
- **ethernet-switching-options redundant-trunk-group group g1 interface**
- **ethernet-switching-options redundant-trunk-group group g1 interface**
- **ethernet-switching-options bpdu-block interface**
- **poe interface vlans pro-bng-mc1-bsd1 interface**

Expanding Interface Range Member and Member Range Statements

All **member** and **member-range** statements in an interface range definition are expanded to generate the final list of interface names for the specified interface range.

Example: Expanding Interface Range Member and Member Range Statements

```
[edit]
  interfaces {
  interface-range range-1 {
    member-range ge-0/0/0 to ge-4/0/20;
    member ge-10/1/1;
    member ge-5/[0-5]/*;
    /*Common configuration is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
}
```

For the **member-range** statement, all possible interfaces between **start-range** and **end-range** are considered in expanding the members. For example, the following **member-range** statement:

member-range ge-0/0/0 to ge-4/0/20

expands to:

```
[ge-0/0/0, ge-0/0/1 ... ge-0/0/max_ports
 ge-0/1/0  ge-0/1/1 ... ge-0/1/max_ports
 ge-0/2/0  ge-0/2/1 ... ge-0/2/max_ports
                     .
                     .
 ge-0/MAX_PICS/0 ... ge-0/max_pics/max_ports
 ge-1/0/0  ge-1/0/1 ... ge-1/0/max_ports
                     .
 ge-1/MAX_PICS/0 ... ge-1/max_pics/max_ports
                     .
                     .
 ge-4/0/0 ge-4/0/1  ... ge-4/0/max_ports]
```


The following **member** statement:

ge-5/[0-5]/*
expands to:

```
ge-5/0/0 ... ge-5/0/max_ports
ge-5/1/0 ... ge-5/0/max_ports
    .
     .
ge-5/5/0 ... ge-5/5/max_ports
```

The following **member** statement:

ge-5/1/[2,3,6,10]
expands to:

```
ge-5/1/2
ge-5/1/3
ge-5/1/6
ge-5/1/10
```

Configuration Inheritance for Member Interfaces

When the Junos OS expands the **member** and **member-range** statements present in an **interface-range**, it creates *interface objects* if they are not explicitly defined in the configuration. The common configuration is copied to all its member interfaces in the **interface-range**.

Example: Configuration Priorities

Foreground interface configuration takes priority compared to configuration inherited by the interface through the **interface-range**.

```
interfaces {
  interface-range range-1 {
    member-range ge-1/0/0/ to ge-10/0/47;
    mtu 256;
    }
    ge-1/0/1 {
      mtu 1024;
    }
}
```

In the preceding example, interface **ge-1/0/1** will have an MTU value of 1024.

This can be verified with output of the **show interfaces | display inheritance** command, as follows:

```
user@host: # show interfaces | display inheritance
## 'ge-1/0/0' was expanded from interface-range 'range-1'
##
ge-1/0/0 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
}
ge-1/0/1 {
    mtu 1024;
```

```
}
##
## 'ge-1/0/2' was expanded from interface-range 'range-1'
##
ge-1/0/2 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
}
      .........
      .........
##
## 'ge-10/0/47' was expanded from interface-range 'range-1'
##
ge-10/0/47 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
}
```

Member Interfaces Inheriting Configuration from Configuration Groups

Interface range member interfaces inherit the config-groups configuration like any other foreground configuration. **interface-range** is similar to any other foreground configuration statement. The only difference is that the **interface-range** goes through a member interfaces expansion before Junos OS reads this configuration.

```
groups {
  global {
    interfaces {
      <*> {
        hold-time up 10;
      }
    }
  }
  apply-groups [global];
  interfaces {
    interface-range range-1 {
      member-range ge-1/0/0 to ge-10/0/47;
      mtu 256;
    }
  }
}
```

The **hold-time** configuration is applied to all members of **interface-range *range-1***.

This can be verified with **show interfaces | display inheritance** as follows:

```
user@host# show interfaces | display inheritance
ge-1/0/0 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
```



```
    ##
    hold-time up 10;
}
ge-1/0/1 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}
ge-10/0/47 {
    ##
    ## '256' was expanded from interface-range 'range-1'
    ##
    mtu 256;
    ##
    ## 'hold-time' was inherited from group 'global'
    ## '10' was inherited from group 'global'
    ##
    hold-time up 10;
}
```

Interfaces Inheriting Common Configuration

If an interface is a member of several interface ranges, that interface will inherit the common configuration from all of those interface ranges.

```
[edit]
interfaces {
  interface-range range-1 {
    member-range ge-1/0/0 to ge-10/0/47;
    mtu 256;
  }
  }
  interfaces {
    interface-range range-1 {
      member-range ge-10/0/0 to ge-10/0/47;
      hold-time up 10;
    }
  }
```

In this example, interfaces **ge-10/0/0** through **ge-10/0/47** will have both **hold-time** and **mtu**.

Configuring Inheritance Range Priorities

The interface ranges are defined in the order of inheritance priority, with the first interface range configuration data taking priority over subsequent interface ranges.

```
[edit]
interfaces {
  interface-range int-grp-one {
    member-range ge-0/0/0 to ge-4/0/40;
    member ge-1/1/1;
```

```
    /*Common config is added part of the interface-range definition*/
    mtu 256;
    hold-time up 10;
  }
  }
  interfaces {
    interface-range int-grp-two {
      member-range ge-5/0/0 to ge-10/0/40;
      member ge-1/1/1;
      mtu 1024;
    }
  }
```

Interface **ge-1/1/1** exists in both **interface-range** ***int-grp-one*** and **interface-range** ***int-grp-two***. This interface inherits **mtu** ***256*** from **interface-range** ***int-grp-one*** because it was defined first.

Configuration Expansion Where Interface Range Is Used

In this example, **interface-range** ***range-1*** is used under the **protocols** hierarchy:

```
[edit]
interfaces {
  interface-range range-1 {
    member ge-10/1/1;
    member ge-5/5/1;
    mtu 256;
    hold-time up 10;
    ether-options {
      flow-control;
      speed {
        100m;
      }
      802.3ad primary;
    }
  }
  protocols {
    dot1x {
      authenticator {
        interface range-1 {
          retries 1;
        }
      }
    }
  }
}
```

The **interface** node present under **authenticator** is expanded into member interfaces of the **interface-range** ***range-1*** as follows:

```
protocols {
  dot1x {
    authenticator {
      interface ge-10/1/1 {
        retries 1;
      }
      interface ge-5/5/1 {
```

```
            retries 1;
          }
        }
      }
    }
```

The **interface** ***range-1*** statement is expanded into two interfaces, ge-10/1/1 and ge-5/5/1, and configuration **retries** ***1*** is copied under those two interfaces.

This configuration can be verified using the **show protocols dot1x | display inheritance** command.

Related Documentation

- *Physical Interfaces*

Configuring Accounting for the Physical Interface

Juniper Networks routers and switches can collect various kinds of data about traffic passing through the router and switch. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management Administration Guide for Routing Devices*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter** ***filter-name*****]** and **[edit firewall family** ***family*** **filter** ***filter-name*****]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Applying an Accounting Profile to the Physical Interface

To enable accounting on an interface, include the **accounting-profile** statement at the **[edit interfaces** ***interface-name*****]** hierarchy level:

```
[edit interfaces interface-name]
accounting-profile name;
```

You can also reference profiles by logical unit; for more information, see "Configuring Accounting for the Logical Interface" on page 114.

Example: Applying an Accounting Profile to the Physical Interface

Configure an accounting profile for an interface and apply it to a physical interface:

```
[edit]
accounting-options {
  file if_stats {
    size 4m files 10 transfer-interval 15;
    archive-sites {
      "ftp://login:password@host/path";
    }
  }
  interface-profile if_profile {
    interval 15;
    file if_stats {
      fields {
        input-bytes;
        output-bytes;
        input-packets;
        output-packets;
        input-errors;
        output-errors;
      }
    }
  }
}
[edit interfaces ge-1/0/1]
accounting-profile if_profile;
```

Configuring Accounting for the Logical Interface

Juniper Networks routers or switches can collect various kinds of data about traffic passing through the router or switch . You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router or switch retains before discarding, and the number of bytes per file
- The period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the **[edit accounting-options]** hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the **[edit accounting-options]** hierarchy level. You configure filter profiles by including the **filter-profile** statement at the **[edit accounting-options]** hierarchy level. For more information, see the *Network Management Administration Guide for Routing Devices*.

You apply filter profiles by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** and **[edit firewall family *family* filter *filter-name*]** hierarchy levels. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Applying an Accounting Profile to the Logical Interface

To enable accounting on a logical interface, include the **accounting-profile** statement:

```
accounting-profile name;
```

You can include this statement at the following hierarchy level:

- **[edit interfaces *interface-name* unit *logical-unit-number*]**

You can also reference profiles for the physical interface; for more information, see "Configuring Accounting for the Physical Interface" on page 113.

Example: Applying an Accounting Profile to the Logical Interface

Configure an accounting profile for an interface and apply it to a logical interface:

```
[edit]
accounting-options {
  file if_stats {
    size 4m files 10 transfer-interval 15;
    archive-sites {
      "ftp://login:password@host/path";
    }
  }
  interface-profile if_profile {
    interval 15;
    file if_stats {
      fields {
        input-bytes;
        output-bytes;
        input-packets;
        output-packets;
        input-errors;
        output-errors;
      }
    }
  }
}
[edit interfaces ge-1/0/1 unit 1]
accounting-profile if_profile;
```

To reference profiles by physical interface, see "Applying an Accounting Profile to the Physical Interface" on page 113. For information about configuring a firewall filter accounting profile, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Configuring Ethernet Loopback Capability

By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the **loopback** statement:

```
loopback;
```

NOTE: If you configure a local loopback on a 1-port 10-Gigabit IQ2 and IQ2-E PIC using the **loopback** statement at the **[edit interfaces *interface-name* gigether-options]** hierarchy level, the transmit-path stops working, causing the remote end to detect a link down.

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

```
[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-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]**
- **[edit interfaces *interface-name* fastether-options]**
- **[edit interfaces *interface-name* gigether-options]**

Related Documentation

- loopback on page 239
- *Ethernet Interfaces Overview*
- EX Series Switches Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests provide duplicate IP address detection. A gratuitous ARP request is a broadcast request for a router's own IP address. If a router or switch sends an ARP request for its own IP address and no ARP replies are received, the router- or switch-assigned IP address is not being used by other nodes. If a router or switch sends an ARP request for its own IP address and an ARP reply is received, the router- or switch-assigned IP address is already being used by another node.

By default, the router or switch responds to gratuitous ARP requests. On Ethernet interfaces, you can disable responses to gratuitous ARP requests. To disable responses to gratuitous ARP requests, include the **no-gratuitous-arp-request** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
no-gratuitous-arp-request;
```

To return to the default—that is, to respond to gratuitous ARP requests—delete the **no-gratuitous-arp-request** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name no-gratuitous-arp-request
```

Gratuitous ARP replies are reply packets sent to the broadcast MAC address with the target IP address set to be the same as the sender's IP address. When the router or switch receives a gratuitous ARP reply, the router or switch can insert an entry for that reply in the ARP cache.

By default, updating the ARP cache on gratuitous ARP replies is disabled on the router or switch. On Ethernet interfaces, you can enable handling of gratuitous ARP replies on a specific interface by including the **gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
gratuitous-arp-reply;
```

To restore the default behavior, include the **no-gratuitous-arp-reply** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
no-gratuitous-arp-reply;
```

Related Documentation

- gratuitous-arp-reply on page 214
- *no-gratuitous-arp-request*
- *Ethernet Interfaces Overview*
- EX Series Switches Interfaces Overview on page 3
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Static ARP Table Entries

To configure static ARP table entries, include the **arp** statement:

```
arp ip-address (mac | multicast-mac) mac-address <publish>;
```

You can include this statement at the following hierarchy levels:

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

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, include the **multicast-mac** statement.

Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***; for example, **0011.2233.4455** or **00:11:22:33:44:55**.

For unicast MAC addresses only, if you include the **publish** option, the router or switch replies to proxy ARP requests.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. By including the **arp** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet policer]** hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the **unnumbered-address** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level.

NOTE: The Junos OS supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

Example: Configuring Static ARP Table Entries

Configure two static ARP table entries on the router or switch's management interface:

```
[edit interfaces]
fxp0 {
  unit 0 {
    family inet {
      address 10.10.0.11/24 {
        arp 10.10.0.99 mac 0001.0002.0003;
        arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
      }
    }
  }
}
```

Related Documentation

- *Management Ethernet Interface Overview*
- EX Series Switches Interfaces Overview on page 3
- *Applying Policers*
- *Configuring an Unnumbered Interface*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Disabling the Transmission of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the **no-redirects** statement:

```
no-redirects;
```


You can include this statement 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*]**

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the **[edit system]** hierarchy level.

Related Documentation

- no-redirects on page 253

Configuring Restricted and Unrestricted Proxy ARP

To configure restricted or unrestricted proxy ARP, include the **proxy-arp** statement:

```
proxy-arp (restricted |unrestricted);
```

You can include this statement at the following hierarchy levels:

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

To return to the default—that is, to disable restricted or unrestricted proxy ARP—delete the **proxy-arp** statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of restricted or unrestricted proxy ARP requests processed by the router or switch by issuing the **show system statistics arp** operational mode command.

NOTE: When proxy ARP is enabled as default or unrestricted, the router or switch responds to any ARP request as long as the device has an active route to the target address of the ARP request. This gratuitous ARP behavior can result in an error when the receiving interface and target response interface are the same and the end device (for example, a client) performs a duplicate address check. To prevent this error, configure the router or switch interface with the **no-gratuitous-arp-reply** statement. See "Configuring Gratuitous ARP" on page 116 for information about how to disable responses to gratuitous ARP requests.

Related Documentation

- proxy-arp on page 263
- *Restricted and Unrestricted Proxy ARP Overview*
- Configuring Gratuitous ARP on page 116
- *Ethernet Interfaces Feature Guide for Routing Devices*

Enabling or Disabling SNMP Notifications on Logical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the logical interface, include the **traps** statement; to disable these notifications on the logical interface, include the **no-traps** statement:

```
(traps | no-traps);
```

You can include these statements at the following hierarchy levels:

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

NOTE: Gigabit Ethernet interfaces on J Series routers do not support SNMP.

Enabling or Disabling SNMP Notifications on Physical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. You can enable or disable these notification based on you requirements.

To explicitly enable sending SNMP notifications on the physical interface, perform the following steps:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. Configure the **traps** option to enable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

   ```
   [edit interfaces interface-name]
   user@host# set traps
   ```

To disable sending SNMP notifications on the physical interface, perform the following steps:

1. In configuration mode, go to the **[edit interfaces *interface-name*]** hierarchy level:

   ```
   [edit]
   user@host# edit interfaces interface-name
   ```

2. Configure the **no-traps** option to disable sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

   ```
   [edit interfaces interface-name]
   user@host# set no-traps
   ```


NOTE: Gigabit Ethernet interfaces on J Series routers do not support SNMP.

Related Documentation

- traps on page 274

Configuring Aggregated Ethernet Links (CLI Procedure)

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

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

To configure aggregated Ethernet interfaces, using the CLI:

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

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

2. Specify the minimum number of links for the aggregated Ethernet interface (ae*x*), 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 ae0 aggregated-ether-options minimum-links number
   ```

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

   ```
   [edit interfaces]
   user@switch# set ae0 aggregated-ether-options link-speed speed
   ```

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

   ```
   [edit interfaces]
   user@switch# set xe-fpc/pic/port ether-options 802.3ad ae0
   user@switch# set xe-fpc/pic/port ether-options 802.3ad ae0
   ```

5. Specify an interface family for the aggregated Ethernet bundle:

   ```
   [edit interfaces]
   user@switch# set ae0 unit 0 family inet address address
   ```

For information about adding LACP to a LAG, see "Configuring Aggregated Ethernet LACP (CLI Procedure)" on page 125.

Related Documentation

- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 122
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125

- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- Verifying the Status of a LAG Interface on page 291
- Understanding Aggregated Ethernet Interfaces and LACP on page 8

Configuring Aggregated Ethernet Interfaces (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

Use the link aggregation feature to aggregate one or more Ethernet interfaces to form a virtual link or link aggregation group (LAG) on an EX Series switch. 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 availability. You can use the J-Web interface to configure aggregated Ethernet interfaces, or a LAG, on the switch.

NOTE: Interfaces that are already configured with MTU, duplex, flow control, or logical interfaces are listed but are not available for aggregation.

To configure an aggregated Ethernet interface (also referred to as a LAG):

1. Select **Configure** > **Interfaces** > **Link Aggregation**.

   The list of aggregated interfaces is displayed.

   NOTE: After you make changes to the configuration on this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See Using the Commit Options to Commit Configuration Changes for details about all commit options.

2. Click one of the following:

   - **Add**—Creates an aggregated Ethernet interface, or LAG. Enter information as specified in Table 30 on page 123.
   - **Edit**—Modifies a selected LAG.

- **Aggregation**—Modifies settings for the selected LAG. Enter information as specified in Table 30 on page 123.
- **VLAN**—Specifies VLAN options for the selected LAG. Enter information as specified in Table 31 on page 124.
- **IP Option**—Specifies IP options for the selected LAG. Enter information as specified in Table 32 on page 124.

- **Delete**—Deletes the selected LAG.
- **Disable Port** or **Enable Port**—Disables or enables the administrative status on the selected interface.
- **Device Count**—Configures the number of aggregated logical devices available to the switch. Select the number and click **OK**.

Table 30: Aggregated Ethernet Interface Options

Field	Function	Your Action
Aggregated Interface	Specifies the name of the aggregated interface.	None. The name is supplied by the software.
LACP Mode	Specifies the mode in which LACP packets are exchanged between the interfaces. The modes are: • **None**—Indicates that no mode is applicable. • **Active**—Indicates that the interface initiates transmission of LACP packets • **Passive**—Indicates that the interface responds only to LACP packets.	Select from the list.
Description	Specifies a description for the LAG.	Enter a description.
Interface	Specifies the interfaces in the LAG.	To add interfaces to the LAG, select the interfaces and click **Add**. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list. Click **OK**. To remove an interface from the LAG, select the interface and click **Remove**. NOTE: Only interfaces that are configured with the same speed can be selected together for a LAG.
Enable Log	Specifies whether to enable generation of log entries for the LAG.	Select the check box to enable log generation, or clear the check box to disable log generation.

Table 31: VLAN Options

Field	Function	Your Action
Port Mode	Specifies the mode of operation for the port: trunk or access.	If you select **Trunk**, you can: 1. Click **Add** to add a VLAN member. 2. Select the VLAN and click **OK**. 3. (Optional) Associate a native VLAN ID with the port. If you select **Access**, you can: 1. Select the VLAN member to be associated with the port. 2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN. Click **OK**.

Table 32: IP Options

Field	Function	Your Action
IPv4 Address	Specifies an IPv4 address for the selected LAG.	1. Select the check box **IPv4 address**. 2. Type an IP address—for example, **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. 4. Click **OK**.
IPv6 Address	Specifies an IPv6 address for the selected LAG.	1. Select the check box **IPv6 address**. 2. Type an IP address—for example, **2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix. 4. Click **OK**.

Related Documentation

- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- Verifying the Status of a LAG Interface on page 291

- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- Understanding Aggregated Ethernet Interfaces and LACP on page 8

Configuring Aggregated Ethernet LACP (CLI Procedure)

For aggregated Ethernet interfaces on EX Series switches, 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 interfaces with or without LACP enabled.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group

NOTE: **You can also configure LACP link protection on aggregated Ethernet interfaces. For information, see "Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)" on page 126.**

The Junos OS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

Before you configure LACP, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See "Configuring Aggregated Ethernet Links (CLI Procedure)" on page 121

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), which contain 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. Configure at least one side of the aggregated Ethernet link as active:

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

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

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

NOTE: The LACP process exists in the system only if you configure the system in either active or passive LACP mode.

Related Documentation

- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 122
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Verifying the Status of a LAG Interface on page 291
- Understanding Aggregated Ethernet Interfaces and LACP on page 8

Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure)

You can configure LACP link protection and system priority at the global level on the switch or for a specific aggregated Ethernet interface. When using LACP link protection to protect a single link in the aggregated ethernet bundle, you configure only two member links for an aggregated Ethernet interface: one active and one standby. LACP link protection ensures that only one link—the link with the higher priority—is used for traffic. The other link is forced to stay in a *waiting* state.

When using LACP link protection to protect multiple links in an aggregated ethernet bundle, you configure links into primary and backup subgroups. A link protection subgroup is a collection of ethernet links within the aggregated ethernet bundle. When you use link protection subgroups, you configure a primary subgroup and a backup subgroup. The configuration process includes assigning member links to each subgroup. When the configuration process is complete, the primary subgroup is used to forward traffic until a switchover event, such as a link failure, occurs and causes the backup subgroup to assume control of traffic that was travelling on the links in the primary subgroup within the bundle.

By default LACP link protection reverts to a higher-priority (lower-numbered) link when the higher-priority link becomes operational or when a higher-priority link is added to the aggregated Ethernet bundle. For priority purposes, LACP link protection treats subgroups like links. You can suppress link calculation by adding the **non-revertive** statement to the link protection configuration. In nonrevertive mode, when a link is active in sending and receiving LACP packets, adding a higher-priority link to the bundle does not change the status of the currently active link. It remains active.

If LACP link configuration is specified to be nonrevertive at the global **[edit chassis]** hierarchy level, you can specify the **revertive** statement in the LACP link protection configuration at the aggregated Ethernet interface level to override the nonrevertive setting for the interface. In revertive mode, adding a higher-priority link to the aggregated Ethernet bundle results in LACP recalculating the priority and switching the status from the currently active link to the newly added, higher-priority link.

NOTE: When LACP link protection is enabled on both local and remote sides of the link, both sides must use the same mode (either revertive or nonrevertive).

Configuring LACP link configuration at the aggregated Ethernet level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

Before you configure LACP link protection, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See “Configuring Aggregated Ethernet Links (CLI Procedure)” on page 121.
- Configured LACP for the interface. See “Configuring Aggregated Ethernet LACP (CLI Procedure)” on page 125.

You can configure LACP link protection for all aggregated Ethernet interfaces on the switch by enabling it at the global level on the switch or configure it for a specific aggregated Ethernet interface by enabling it on that interface.

Configuring LACP Link Protection for a Single Link at the Global Level

To configure LACP link protection for aggregated Ethernet interfaces at the global level:

1. Enable LACP link protection on the switch:

   ```
   [edit chassis aggregated-devices ethernet lacp]
   user@switch# set link-protection
   ```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interfaces to be in nonrevertive mode:

   NOTE: LACP link protection is in revertive mode by default.

   ```
   [edit chassis aggregated-devices ethernet lacp link-protection]
   user@switch# set non-revertive
   ```

3. (Optional)To configure LACP system priority for the aggregated Ethernet interfaces:

   ```
   [edit chassis aggregated-devices ethernet lacp]
   user@switch# set system-priority
   ```

Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level

To enable LACP link protection for a specific aggregated Ethernet interface:

1. Enable LACP link protection for the interface:

   ```
   [edit interfaces aeX aggregated-ether-options lacp]
   user@switch# set link-protection
   ```

2. (Optional) Configure the LACP link protection for the aggregated Ethernet interface to be in revertive or nonrevertive mode:

   - To specify revertive mode:

     ```
     [edit interfaces aeX aggregated-ether-options lacp link-protection]
     user@switch# set revertive
     ```

   - To specify nonrevertive mode:

     ```
     [edit interfaces aeX aggregated-ether-options lacp link-protection]
     user@switch# set non-revertive
     ```

3. (Optional)To configure LACP system priority for an aggregated Ethernet interface:

   ```
   [edit interfaces aeX aggregated-ether-options lacp link-protection]
   user@switch# set system-priority
   ```

4. (Optional) To configure LACP port priority for an aggregated Ethernet interface:

   ```
   [edit interfaces ge-fpc/pic/port ether-options 802.3ad lacp]
   user@switch# set port-priority
   ```

Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface

You can configure link protection subgroup bundles to provide link protection for multiple links in an aggregated ethernet bundle.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a LAG bundle, instead of providing protection to a single link in the aggregated ethernet bundle only. You can, for instance, configure a primary subgroup with three member links and a backup subgroup with three different member links and use the backup subgroup to provide link protection for the primary subgroup.

To configure link protection using subgroups:

1. Configure the primary link protection subgroup in the aggregated ethernet interface:

   ```
   [edit interfaces aeX aggregated-ether-options]
   user@switch# set link-protection-sub-group group-name primary
   ```

   For instance, to create a primary link protection subgroup named **subgroup-primary** for interface **ae0**:

   ```
   [edit interfaces ae0 aggregated-ether-options]
   user@switch# set link-protection-sub-group subgroup-primary primary
   ```

2. Configure the backup link protection subgroup in the aggregated ethernet interface:

   ```
   [edit interfaces aeX aggregated-ether-options]
   user@switch# set link-protection-sub-group group-name backup
   ```

   For instance, to create a backup link protection subgroup named **subgroup-backup** for interface **ae0**:

   ```
   [edit interfaces ae0 aggregated-ether-options]
   user@switch# set link-protection-sub-group subgroup-backup backup
   ```

   NOTE: You can create one primary and one backup link protection subgroup per aggregated ethernet interface.

3. Attach interfaces to the link protection subgroups:

   ```
   [edit interfaces interface-name ether-options 802.3ad]
   user@switch# set link-protection-sub-group group-name
   ```

   NOTE: The primary and backup link protection subgroups must contain the same number of interfaces. For instance, if the primary link protection subgroup contains three interfaces, the backup link protection subgroup must also contain three interfaces.

   For instance, to configure interfaces **ge-0/0/0** and **ge-0/0/1** into link protection subgroup **subgroup-primary** and interfaces **ge-0/0/2** and **ge-0/0/3** into link protection subgroup **subgroup-backup**:

   ```
   [edit interfaces ge-0/0/0 ether-options 802.3ad]
   user@switch# set link-protection-sub-group subgroup-primary
   ```

```
[edit interfaces ge-0/0/1 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-primary
[edit interfaces ge-0/0/2 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
[edit interfaces ge-0/0/3 ether-options 802.3ad]
user@switch# set link-protection-sub-group subgroup-backup
```

4. (Optional) Configure the port priority for link protection:

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

   The port priority is used to select the active link.

5. Enable link protection

   To enable link protection at the LAG level:

   ```
   [edit interfaces aeX aggregated-ether-options]
   user@switch# set link-protection
   ```

   To enable link protection at the LACP level:

   ```
   [edit interfaces aeX aggregated-ether-options lacp]
   user@switch# set link-protection
   ```

   For instance, to enable link protection on **ae0** at the LAG level:

   ```
   [edit interfaces ae0 aggregated-ether-options]
   user@switch# set link-protection
   ```

   For instance, to enable link protection on **ae0** at the LACP level:

   ```
   [edit interfaces ae0 aggregated-ether-options lacp]
   user@switch# set link-protection
   ```

Related Documentation

- Understanding Aggregated Ethernet Interfaces and LACP on page 8

Configuring Aggregated Ethernet Link Protection

You can configure link protection for aggregated Ethernet interfaces to provide QoS on the links during operation.

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

NOTE: Link protection is not supported on MX80.

- Configuring Link Protection for Aggregated Ethernet Interfaces on page 131
- Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces on page 131

Configuring Link Protection for Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces support link protection to ensure QoS on the interface.

To configure link protection:

1. Specify that you want to configure the options for an aggregated Ethernet interface.

   user@host# **edit interfaces ae*x* aggregated-ether-options**

2. Configure the link protection mode.

   [edit interfaces ae*x* aggregated-ether-options]
   user@host# **set** link-protection

Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces

To configure link protection, you must specify a primary and a secondary, or backup, link.

To configure a primary link and a backup link:

1. Configure the primary logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* primary**

2. Configure the backup logical interface.

   [edit interfaces *interface-name*]
   user@host# **set (fastether-options | gigether-options) 802.3ad ae*x* backup**

Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link

On aggregated Ethernet interfaces, you designate a primary and backup link to support link protection. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router or switch. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link.

To manually control when traffic should be diverted back to the primary link from the designated backup link, enter the following operational command:

user@host> **request interface revert ae*x***

Disabling Link Protection for Aggregated Ethernet Interfaces

To disable link protection, issue the **delete interface revert ae*x*** configuration command.

user@host# **delete interfaces ae*x* aggregated-ether-options link-protection**

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. Generally, all interfaces that make up a bundle must have the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed that you specify in the **link-speed** parameter, an error message is logged. However, starting with Junos OS Release 13.2, aggregated Ethernet supports the following mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers:

- Member links of different modes (WAN and LAN) for 10-Gigabit Ethernet links.
- Member links of different rates: 10-Gigabit Ethernet, 40-Gigabit Ethernet, 50-Gigabit Ethernet, 100-Gigabit Ethernet, and OC192 (10-Gigabit Ethernet WAN mode)

NOTE:

- Member links of 50-Gigabit Ethernet can only be configured using the 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP (PD-1CE-CFP-FPC4).
- Starting with Junos OS Release 13.2, 100-Gigabit Ethernet member links can be configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP. This 100-Gigabit Ethernet member link can be included in an aggregated Ethernet link that includes member links of other interfaces as well. In releases before Junos OS Release 13.2, the 100-Gigabit Ethernet member link configured using the two 50-Gigabit Ethernet interfaces of 100-Gigabit Ethernet PIC with CFP cannot be included in an aggregated Ethernet link that includes member links of other interfaces.

To configure member links of mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers, you need to configure the **mixed** option for the [**edit interfaces** ***aex*** **aggregated-ether-options link-speed**] statement.

To set the required link speed:

1. Specify that you want to configure the aggregated Ethernet options.

   user@host# **edit interfaces** ***interface-name*** **aggregated-ether-options**

2. Configure the link speed.

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

   speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

   Aggregated Ethernet interfaces on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.
- **50g**—Links are 50 Gbps.

Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.
- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

Related Documentation

- *aggregated-ether-options*
- *Configuring Mixed Aggregated Ethernet Links*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled **up**. By default, only one link must be up for the bundle to be labeled **up**.

To configure the minimum number of links:

1. Specify that you want to configure the aggregated Ethernet options.

user@host# **edit interfaces *interface-name* aggregated-ether-options**

2. Configure the minimum number of links.

   [edit interfaces *interface-name* aggregated-ether-options]
   user@host# **set** minimum-links ***number***

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, and EX 9200 switches, the valid range for **minimum-links *number*** is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

On EX8200 switches, the range of valid values for **minimum-links *number*** is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled **up**.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the **aggregated-ether-options** statement, the configuration commit fails and an error message is displayed.

Related Documentation

- *aggregated-ether-options*
- minimum-links on page 246
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the **vlan-tagging** statement at the **[edit interfaces ae*x*]** hierarchy level:

[edit interfaces ae*x*]
vlan-tagging;

You must also include the **vlan-id** statement:

vlan-id *number*;

You can include this statement at the following hierarchy levels:

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

For more information about the **vlan-tagging** and **vlan-id** statements, see “802.1Q VLANs Overview” on page 19.

Related Documentation

- *vlan-id*
- vlan-tagging on page 280

Configuring a Layer 3 Subinterface (CLI Procedure)

EX Series switches use Layer 3 subinterfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. The switch uses the Layer 3 subinterfaces to route traffic between subnets.

To configure Layer 3 subinterfaces, you enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See *Configuring VLANs for EX Series Switches (CLI Procedure)* or *Configuring VLANs for EX Series Switches (J-Web Procedure)*.

To configure Layer 3 subinterfaces:

1. Enable VLAN tagging:

```
[edit interfaces interface-name]
user@switch# set vlan-tagging
```

2. Bind each VLAN ID to a logical interface:

```
[edit interfaces interface-name]
user@switch# set unit logical-unit-number vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number
```

Related Documentation

- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 44
- Verifying That Layer 3 Subinterfaces Are Working on page 293
- Understanding Layer 3 Subinterfaces on page 11

Configuring Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Enabling unicast RPF on the switch interfaces filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. When a packet comes into an interface, if that interface is not the best return path to the source, the switch discards the packet. If the incoming interface is the best return path to the source, the switch forwards the packet.

NOTE: **On EX3200, EX4200, and EX4300 switches, you can enable unicast RPF only globally—that is, on all switch interfaces. You cannot enable unicast RPF on a per-interface basis.**

Before you begin:

- On an EX8200, EX6200, QFX Series switch, or OCX Series switch, ensure that the selected switch interface is symmetrically routed before you enable unicast RPF. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

- On an EX3200, EX4200, or EX4300 switch, ensure that *all* switch interfaces are symmetrically routed before you enable unicast RPF on an interface. When you enable unicast RPF on any interface, it is enabled globally on all switch interfaces. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

To enable unicast RPF, configure it explicitly on a selected customer-edge interface:

[edit interfaces]
user@switch# **set ge-1/0/10 unit 0 family inet rpf-check**

BEST PRACTICE: On EX3200, EX4200, and EX4300 switches, unicast RPF is enabled globally on *all* switch interfaces, regardless of whether you configure it explicitly on only one interface or only on some interfaces.

On EX3200, EX4200, and EX4300 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback of this approach is that the switch displays the flag that indicates that unicast RPF is enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.
- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback of this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

Related Documentation

- Example: Configuring Unicast RPF on an EX Series Switch on page 51
- Verifying Unicast RPF Status on page 293
- Disabling Unicast RPF (CLI Procedure) on page 137
- Troubleshooting Unicast RPF on page 435
- Understanding Unicast RPF on page 12

Disabling Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), disable unicast RPF.

To disable unicast RPF on an EX3200, EX4200, or EX4300 switch, you must delete it from every interface on which you explicitly configured it. If you do not disable unicast RPF on every interface on which you explicitly enabled it, it remains implicitly enabled on all interfaces. If you attempt to delete unicast RPF from an interface on which it was not explicitly enabled, the **warning: statement not found** message appears. If you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces of the EX3200, EX4200, or EX4300 switch.

On EX8200, EX6200, QFX Series switches, and OCX Series switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

To disable unicast RPF, delete its configuration from the interface:

[edit interfaces]
user@switch# **delete ge-1/0/10 unit 0 family inet** rpf-check

NOTE: On EX3200, EX4200, and EX4300 switches, if you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces.

Related Documentation

- Example: Configuring Unicast RPF on an EX Series Switch on page 51
- Verifying Unicast RPF Status on page 293
- Configuring Unicast RPF (CLI Procedure) on page 135
- Understanding Unicast RPF on page 12

Configuring IP Directed Broadcast (CLI Procedure)

NOTE: This task uses Junos OS for EX Series switches that does not support the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Configuring IP Directed Broadcast (CLI Procedure)*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

You can use IP directed broadcast on an EX Series switch to facilitate remote network management by sending broadcast packets to hosts on a specified subnet without broadcasting to the entire network. IP directed broadcast packets are broadcast on only the target subnet. The rest of the network treats IP directed broadcast packets as unicast packets and forwards them accordingly.

Before you begin to configure IP directed broadcast:

- Ensure that the subnet on which you want broadcast packets using IP direct broadcast is not directly connected to the Internet.
- Configure a routed VLAN interface (RVI) for the subnet that will be enabled for IP direct broadcast. See *Configuring Routed VLAN Interfaces (CLI Procedure)* or *Configuring VLANs for EX Series Switches (J-Web Procedure)*.

NOTE: **We recommend that you do not enable IP directed broadcast on subnets that have a direct connection to the Internet because of increased exposure to denial-of-service (DoS) attacks.**

To enable IP directed broadcast for a specified subnet:

1. Add the target subnet's logical interfaces to the VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```

2. Configure the Layer 3 interface on the VLAN that is the target of the IP directed broadcast packets:

```
[edit interfaces]
user@switch# set vlan.1 family inet address 10.1.2.1/24
```

3. Associate a Layer 3 interface with the VLAN:

```
[edit vlans]
user@switch# set v1 l3-interface (VLANs) vlan.1
```

4. Enable the Layer 3 interface for the VLAN to receive IP directed broadcasts:

```
[edit interfaces]
user@switch# set vlan.1 family inet targeted-broadcast
```

Related Documentation

- Example: Configuring IP Directed Broadcast on an EX Series Switch on page 56
- Understanding IP Directed Broadcast for EX Series Switches on page 16

Tracing Operations of an Individual Router or Switch Interface

To trace the operations of individual router or switch interfaces, include the **traceoptions** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
traceoptions {
  flag flag;
}
```


You can specify the following interface tracing flags:

- **all**—Trace all interface operations.
- **event**—Trace all interface events.
- **ipc**—Trace all interface interprocess communication (IPC) messages.
- **media**—Trace all interface media changes.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** files.

Related Documentation

- Tracing Operations of the Interface Process on page 139
- *Tracing Interface Operations Overview*

Tracing Operations of the Interface Process

To trace the operations of the router or switch interface process, dcd, perform the following steps:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level:

   [edit]
   user@host# **edit interfaces**

2. Configure the **traceoptions** statement.

   [edit interfaces]
   user@host# **edit traceoptions**

3. Configure the **no-remote-trace** option to disable remote tracing.

   [edit interfaces traceoptions]
   user@host# **set no-remote-trace**

4. Configure the **file *filename*** option.

   [edit interfaces traceoptions]
   user@host# **edit file**

5. Configure the **files *number*** option, **match *regular-expression*** option, **size *size*** option, and **world-readable | no-world-readable** option.

   [edit interfaces traceoptions file]
   user@host# **set files *number***
   user@host# **set match *regular-expression***
   user@host# **set size *size***
   user@host# **set word-readable | no-world-readable**

6. Configure the tracing flag.

   [edit interfaces traceoptions]
   user@host# **set flag *flag-option***

7. Configure the **disable** option in **flag *flag-option*** statement to disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.


```
[edit interfaces traceoptions]
user@host# set flag flag-option disable
```

You can specify the following flags in the **interfaces traceoptions** statement:

- **all**—Enable all configuration logging.
- **change-events**—Log changes that produce configuration events.
- **gres-events**—Log the events related to GRES.
- **resource-usage**—Log the resource usage for different states.
- **config-states**—Log the configuration state machine changes.
- **kernel**—Log configuration IPC messages to kernel.
- **kernel-detail**—Log details of configuration messages to kernel.
- **select-events**—Log the events on select state machine.

By default, interface process operations are placed in the file named dcd and three 1-MB files of tracing information are maintained.

For general information about tracing, see the tracing and logging information in the *Junos OS Administration Library for Routing Devices*.

Related Documentation

- *Tracing Interface Operations Overview*
- *Tracing Operations of an Individual Router Interface*
- traceoptions on page 273

Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)

SFP+ uplink modules are supported on EX3200 and EX4200 switches, and SFP+ Media Access Control Security (MACSec) uplink modules are supported on EX4200 switches. You can use these uplink modules either for two SFP+ transceivers or four SFP transceivers. You configure the operating mode on the module to match the type of transceiver you want to use—that is, for SFP+ transceivers, you configure the 10-gigabit operating mode, and for SFP transceivers, you configure the 1-gigabit operating mode.

By default, the SFP+ uplink module operates in the 10-gigabit mode and supports only SFP+ transceivers. If you have not changed the module from the default setting and you want to use SFP+ transceivers, you do not need to configure the operating mode.

To set the operating mode of an SFP+ or SFP+ MACSec uplink module:

1. Change the operating mode to the appropriate mode for the transceiver type you want to use by using one of the following commands:

   ```
   [edit]
   user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 1g
   [edit]
   user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 10g
   ```
2. (SFP+ uplink module only) If the switch is running:

- Junos OS Release 10.1 or later, the changed operating mode takes effect immediately unless a port on the SFP+ uplink module is a Virtual Chassis port (VCP). If any port on the SFP+ uplink module is a VCP, the changed operating mode does not take effect until the next reboot of the switch.

  NOTE: **During the operating mode change, the Packet Forwarding Engine is restarted. In a Virtual Chassis configuration, this means that the Flexible PIC Concentrator connection with the master is dropped and then reconnected.**

- Junos OS Release 10.0 or earlier, reboot the switch.

You can see whether the operating mode has been changed to the new mode you configured by issuing the **show chassis pic fpc-slot *slot-number* pic-slot 1** command.

Related Documentation

- *Uplink Modules in EX3200 Switches*
- *Uplink Modules in EX4200 Switches*
- *Pluggable Transceivers Supported on EX3200 Switches*
- *Pluggable Transceivers Supported on EX4200 Switches*

Configuring the Media Type on Dual-Purpose Uplink Ports (CLI Procedure)

EX2200-C switches and ACX1000 routers provide two dual-purpose uplink ports. Each dual uplink port is a single interface that offers a choice of two connections: an RJ-45 connection for a copper Ethernet cable and an SFP connection for a fiber-optic Ethernet cable. You can choose to use either connection, but only one connection can be active at a time.

By default, if you plug a transceiver into the SFP connector, the port becomes a fiber-optic Gigabit Ethernet port, even if a copper Ethernet cable is plugged into the RJ-45 connection as well. If a transceiver is not plugged into the SFP connector, the port defaults to a copper 10/100/1000 Ethernet port.

You can constrain the use of the port to one connection type by configuring the media type for the port to be either copper or fiber. When you configure a media type on the port, the port will no longer accept the alternate connection type. For example, if you configure the uplink port as a fiber port and then plug a copper Ethernet cable into the RJ-45 connector, the interface will not come up.

To configure the media type for an uplink port:

`user@switch#` **set interfaces *interface-name*** media-type (Dual-Purpose Uplink Ports) ***media-type***

For example, to set the media type for uplink port **ge-0/1/0** to copper:

`user@switch#` **set interfaces ge-0/1/0 media-type copper**

NOTE: When you change the media type setting for a dual-purpose uplink port, it can take up to 6 seconds for the interface to appear in operational commands.

Related Documentation

- *EX2200 Switches Hardware Overview*

Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)

Tunneling provides a private, secure path for transporting packets through an otherwise public network by encapsulating packets inside a transport protocol known as an *IP encapsulation protocol*. Generic routing encapsulation (GRE) is an IP encapsulation protocol that is used to transport packets over a network. Information is sent from one network to the other through a GRE tunnel.

GRE tunneling is accomplished through routable tunnel endpoints that operate on top of existing physical and other logical endpoints. GRE tunnels connect one endpoint to another and provide a clear data path between them.

This topic describes:

Configuring a GRE Tunnel Port

To configure GRE tunnels on a switch, you convert a network port or uplink port on the switch to a GRE tunnel port for tunnel services. Each physical tunnel port, named gr-*fpc/pic/port*, can have one or more logical interfaces, each of which is a GRE tunnel.

After conversion to a GRE tunnel port, the physical port cannot be used for network traffic.

To configure a GRE tunnel port on a switch:

1. Determine the network port or uplink port on your switch to convert to a GRE tunnel port.
2. Configure the port as a tunnel port for GRE tunnel services:

```
[edit chassis]
user@switch# set fpc slot pic pic-number tunnel-port port-number tunnel-services
```


NOTE: On EX3200 switches and standalone EX4200, EX4500, and EX4550 switches, the FPC number is 0 because it refers to the switch itself. On EX4200, EX4500, and EX4550 Virtual Chassis, the FPC number is the member ID of the Virtual Chassis member on which you are configuring the port. On EX6200 and EX8200 switches, the FPC number is the number of the slot containing the line card on which you are configuring the port.

For built-in ports on EX3200, EX4200, EX4500, and EX4550 switches and on EX6200 and EX8200 switches, the PIC number is 0. For uplink ports on EX3200 and EX4200 switches, the PIC number is 1.

Configuring Tunnels to Use Generic Routing Encapsulation

Normally, a GRE tunnel port comes up as soon as it is configured and stays up as long as a valid tunnel source address exists or an interface is operational. Each logical interface you configure on the port can be configured as the source or as the endpoint of a GRE tunnel.

To configure a tunnel port to use GRE:

1. Configure a physical GRE port with a logical interface name and address:
   - For IPv4 over GRE, specify the protocol family **inet**:

     [edit interfaces]
     user@switch# **set gr-*fpc/pic/port* unit *number* family inet *address***
   - For IPv6 over GRE, specify the protocol family **inet6**:

     [edit interfaces]
     user@switch# **set gr-*fpc/pic/port* unit *number* family inet6 *address***
2. Specify the tunnel source address for the logical interface:

   [edit interfaces]
   user@switch# **set gr-*fpc/pic/port* unit *number* tunnel source *source-address***
3. Specify the destination address:

   [edit interfaces]
   user@switch# **set gr-*fpc/pic/port* unit *number* tunnel destination *destination-address***

Related Documentation

- Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 296
- Understanding Generic Routing Encapsulation on page 20

Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure)

EX8200 switches support multicast load balancing on link aggregation groups (LAGs). Multicast load balancing evenly distributes Layer 3 multicast routed traffic over the aggregated links. You can aggregate up to twelve 10-gigabit Ethernet links to form a virtual link or a LAG. The MAC client can treat this virtual link as if it were a single link to increase bandwidth, provide graceful degradation if a failure occurs, and increase availability. On EX8200 switches, multicast load balancing is enabled by default. However, if it is explicitly disabled, you can reeable it.

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

To configure a LAG and ensure that multicast load balancing is enabled:

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

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

2. Specify the minimum number of links for the aggregated Ethernet interface (ae*x*)—that is, the defined LAG—to be labeled **up**:

   NOTE: By default, only one link needs to be up for the LAG to be labeled up.

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

3. Specify each member to be included in the LAG:

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

4. To reenable multicast load balancing:

   ```
   [edit chassis]
   user@switch#  set multicast-loadbalance
   ```

   NOTE: You can disable multicast load balancing on an EX8200 switch by using the **set chassis multicast-loadbalance disable** command.

   NOTE: If the EX8200 switch is already handling multicast streams, then you reenable multicast load balancing, the existing steams might not be distributed evenly. In this case, we recommend clearing data related to these streams by using the **clear pim join all** and **clear igmp-snooping membership** commands.

The default multicast load-balancing hashing algorithm, crc-sgip, involves the cyclic redundancy check of source and group IP addresses. You can optionally change the value of the **hash-mode** option in the multicast-loadbalance statement to try different algorithms until you find the one that best distributes your Layer 3 routed multicast traffic.

If you change the hashing algorithm when multicast load balancing is disabled, the new algorithm takes effect after you reenable multicast load balancing.

NOTE: **You do not need to set link speed, which is automatically determined by the size of the LAG that uses multicast load balancing. For example, the link speed of a 40-gigabit LAG would be 40 gigabits per second.**

Related Documentation

- Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches on page 59
- Understanding Multicast Load Balancing on Aggregated 10-Gigabit Links for Routed Multicast Traffic on EX8200 Switches on page 24

Configuring Energy Efficient Ethernet on Interfaces (CLI Procedure)

Energy Efficient Ethernet (EEE), an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, reduces the power consumption of physical layer devices (PHYs) during periods of low link utilization. EEE saves energy by putting part of the transmission circuit into low power mode when a link is idle.

NOTE: **Configure EEE only on EEE-capable Base-T copper Ethernet ports. If you configure EEE on unsupported ports, the console displays the message: “EEE not supported”.**

This topic describes:

- Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port on page 145
- Disabling EEE on a Base-T Copper Ethernet Port on page 145

Enabling EEE on an EEE-Capable Base-T Copper Ethernet Port

To enable EEE on an EEE-capable Base-T copper Ethernet interface:

```
[edit]
user@switch# set interfaces interface-name ether-options ieee-802-3az-eee
```

You can view the EEE status by using the **show interfaces** ***interface-name*** **detail** command.

Disabling EEE on a Base-T Copper Ethernet Port

To disable EEE on a Base-T copper Ethernet interface:

```
[edit]
user@switch# delete interfaces interface-name ether-options ieee-802-3az-eee
```

By default, EEE is disabled on EEE-capable ports.

Related Documentation

- Verifying That EEE Is Saving Energy on Configured Ports on page 296
- Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 27

Damping Shorter Physical Interface Transitions

By default, when an interface changes from being up to being down, or from down to up, this transition is advertised immediately to the hardware and Junos OS. In some situations—for example, when an interface is connected to an add/drop multiplexer (ADM) or wavelength-division multiplexer (WDM), or to protect against SONET/SDH framer holes—you might want to damp interface transitions. This means not advertising the interface's transition until a certain period of time has passed, called the *hold-time*. When you have damped interface transitions and the interface goes from up to down, the down hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *down*, then the router begins to advertise the interface as being down. Similarly, when an interface goes from down to up, the up hold-time timer is triggered. Every interface transition that occurs during the hold-time is ignored. When the timer expires and the interface state is still *up*, then the router begins to advertise the interface as being up. For information about physical interface damping, see *Physical Interface Damping Overview*.

This task applies to damping shorter physical interface transitions in milliseconds. To damp longer physical interface transitions in seconds, see *Damping Longer Physical Interface Transitions*.

To configure damping of shorter physical interface transitions:

1. Select the interface to damp, where the interface name is ***interface-type-fpc/pic/port***:

   [edit]
   user@host# **edit interfaces *interface-name***

2. Configure the hold-time for link up and link down.

   [edit interfaces *interface-name*]
   user@host# **set hold-time up *milliseconds* down *milliseconds***

The hold time can be a value from 0 through 4,294,967,295 milliseconds. The default value is 0, which means that interface transitions are not damped. Junos OS advertises the transition within 100 milliseconds of the time value you specify.

For most Ethernet interfaces, hold timers are implemented using a one-second polling algorithm. For 1-port, 2-port, and 4-port Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs), hold timers are interrupt-driven.

NOTE: **The hold-time option is not available for controller interfaces.**

Related Documentation

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

- *SONET/SDH Defect Hold Times Overview*
- *Configuring SONET/SDH Defect Triggers*
- **hold-time on page 215**

Configuring Local Link Bias (CLI Procedure)

Local link bias is used to conserve bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a link aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF on a different member link in the LAG bundle.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on the LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced across the member links in the LAG bundle as it exits the Virtual Chassis or VCF.

Local link bias can be enabled or disabled globally or per LAG bundle on a Virtual Chassis or VCF. In cases where local link bias is enabled at both the global and per LAG bundle levels, the per LAG bundle configuration takes precedence. For instance, if local link bias is enabled globally but disabled on a LAG bundle named **ae1**, local link bias is disabled on the LAG bundle named **ae1**.

Enabling Local Link Bias Globally

When local link bias is enabled globally, local link bias is enabled on all LAG bundles in the Virtual Chassis or VCF except the LAG bundles that have explicitly disabled local link bias.

To enable local link bias globally in a Virtual Chassis or VCF:

```
[edit]
user@switch# set forwarding-options local-bias (forwarding-options)
```

Enabling Local Link Bias on a Single LAG Bundle in a Virtual Chassis or Virtual Chassis Fabric

When local link bias is enabled for a single LAG bundle, it remains enabled for that LAG bundle regardless of the global local link bias setting. If local link bias is enabled on a single LAG bundle but disabled globally, for instance, local link bias is enabled for the

LAG bundle that was individually configured to enable local link bias but disabled for all other LAG bundles in the Virtual Chassis or VCF.

To enable local link bias on an individual LAG bundle:

```
[edit]
user@switch# set interface aex aggregated-ether-options local-bias (edit interfaces ae)
```

where **ae*x*** is the name of the aggregated Ethernet link bundle.

For instance, to enable local link bias on aggregated Ethernet interface ae0:

```
[edit]
user@switch# set interface ae0 aggregated-ether-options local-bias
```

Disabling Local Link Bias Globally in a Virtual Chassis or Virtual Chassis Fabric

You can disable local link bias globally if it was previously globally enabled.

When local link bias is disabled globally, local link bias is disabled on all LAG bundles in the Virtual Chassis or VCF except the LAG bundles that have explicitly enabled local link bias.

To disable local link bias globally in a Virtual Chassis or VCF:

```
[edit]
user@switch# set forwarding-options local-bias (forwarding-options) disable
```

Disabling Local Link Bias on a Single LAG Bundle in a Virtual Chassis or Virtual Chassis Fabric

When local link bias is disabled for a single LAG bundle, it remains disabled for that LAG bundle regardless of the global local link bias setting. If local link bias is disabled on the LAG bundle named **ae1**, for instance, but enabled globally, local link bias is disabled on the LAG bundle named **ae1**.

To disable local link bias on an individual LAG bundle:

```
[edit]
user@switch# set interface aex aggregated-ether-options local-bias (edit interfaces ae) disable
```

where **ae*x*** is the name of the aggregated Ethernet link bundle.

For instance, to disable local link bias on aggregated Ethernet interface ae0:

```
[edit]
user@switch# set interface ae0 aggregated-ether-options local-bias disable
```

Related Documentation

- Understanding Local Link Bias on page 28

CHAPTER 4

Configuration Statements

[edit chassis] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit chassis]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

Supported Statements in the [edit chassis] Hierarchy Level

The following hierarchy shows the **[edit chassis]** configuration statements supported on EX Series switches:

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
      lacp {
        link-protection non-revertive;
        system-priority system-priority-number
      }
    }
  }
  alarm {
    ethernet {
      link-down (ignore | red | yellow);
    }
    management-ethernet {
      link-down (ignore | red | yellow);
    }
  container-devices {
    device-count device-count-number;
  }
  disk-partition {
    /config {
      level (full | high) {
        free-space (free-space-threshold-value | mb | percent);
        }
      }
    }
    /var {
      level (full | high) {
        free-space (free-space-threshold-value | mb | percent);
        }
      }
    }
  }
  fpc slot-number{
    pic pic-number {
      no-multi-rate;
      q-pic-large-buffer (large-scale | small-scale);
```



```
            }
         }
      }
      maximum-ecmp maximum-ecmp-routes;
      lcd-menu {
         fpc slot-number {
            menu-item menu-name);
               disable;
      pseudowire-service {
         device-count device-count-number;
      }
      psu {
         redundancy {
            n-plus-n;
         }
      redundancy {
         graceful-switchover;
      }
      slow-pfe-alarm;
   }
```

Related Documentation

- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- *Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)*
- *Configuring Power Supply Redundancy (CLI Procedure)*
- *Configuring the Power Priority of Line Cards (CLI Procedure)*
- *Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)*

[edit interfaces] Configuration Statement Hierarchy on EX Series Switches

Each of the following topics lists the statements at a subhierarchy of the **[edit interfaces]** hierarchy:

- [edit interfaces ae] Configuration Statement Hierarchy on EX Series Switches on page 154
- [edit interfaces ge] Configuration Statement Hierarchy on EX Series Switches on page 158
- *[edit interfaces gr] Configuration Statement Hierarchy on EX Series Switches*
- [edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches on page 161
- [edit interfaces lo] Configuration Statement Hierarchy on EX Series Switches on page 168
- [edit interfaces me] Configuration Statement Hierarchy on EX Series Switches on page 171
- [edit interfaces vlan] Configuration Statement Hierarchy on EX Series Switches on page 175

- [edit interfaces vme] Configuration Statement Hierarchy on EX Series Switches on page 178
- [edit interfaces xe] Configuration Statement Hierarchy on EX Series Switches on page 181

Related Documentation

- EX Series Switches Interfaces Overview on page 3
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 135
- *Configuring Routed VLAN Interfaces (CLI Procedure)*
- *Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)*
- *Junos OS Interfaces Fundamentals Configuration Guide*
- *Junos OS Ethernet Interfaces Configuration Guide*

[edit interfaces ae] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces ae]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

- Supported Statements in the [edit interfaces ae] Hierarchy Level on page 154
- Unsupported Statements in the [edit interfaces ae] Hierarchy Level on page 157

Supported Statements in the [edit interfaces ae] Hierarchy Level

The following hierarchy shows the **[edit interfaces ae]** configuration statements supported on EX Series switches.

```
interfaces {
  ae-fpc/pic/port {
    accounting-profile name;
    aggregated-ether-options {
      ethernet-switch-profile {
        tag-protocol-id identifier ;
      }
      (flow-control | no-flow-control);
```

```
    lacp {
      (active | passive);
      admin-key key;
      periodic interval;
      system-id mac-address;
    }
    (link-protection | no-link-protection);
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
  }
  description text;
  disable;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    arp-resp;
    bandwidth rate;
    description text;
    disable;
    family ccc;
    family ethernet-switching {
      filter {
        input filter-name;
        output filter-name;
      }
      native-vlan-id vlan-id-number;
      port-mode (access | trunk);
      vlan {
        members [ members];
      }
    }
    family inet {
      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;
          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 ];
            virtual-link-local-address address;
            vrrp-inherit-from {
                active-group group-number;
                active-interface interface-name;
            }
        }
    }
    dhcp {
        client-identifier (ascii client-id | hexadecimal client-id);
        lease-time (seconds | infinte);
        retransmission-attempt  number;
        retransmission-interval sections;
        server-address ip-address;
        update-server server;
        vendor-id id;
    }
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    primary;
    rpf-check;
    targeted-broadcast;
}
family inet6 {
    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;
            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 ( address | routing-instance routing-instance-name );
              }
              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 {
            input filter-name;
            output filter-name;
          }
          mtu bytes;
          no-neighbor-learn;
          policer {
            input policer-name;
            output policer-name;
          }
          rpf-check;
        }
        family iso {
          address interface-address;
          mtu bytes;
        }
        family mpls {
          mtu bytes;
        }
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
      }
      vlan-tagging;
    }
  }
```

Unsupported Statements in the [edit interfaces ae] Hierarchy Level

All statements in the **[edit interfaces ae]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 33: Unsupported [edit interfaces ae] Configuration Statements on EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as *interface-range*, are not shown in the statements or hierarchies.	
family fibre-channel	[edit interfaces ae unit]
source-address-filter	[edit interfaces ae aggregated-ether-options]
source-address-filtering | no-source-address-filtering	[edit interfaces ae aggregated-ether-options]

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153

[edit interfaces ge] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces ge]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

- Supported Statements in the [edit interfaces ge] Hierarchy Level on page 158
- Unsupported Statements in the [edit interfaces ge] Hierarchy Level on page 161

Supported Statements in the [edit interfaces ge] Hierarchy Level

The following hierarchy shows the **[edit interfaces ge]** configuration statements supported on EX Series switches.

```
interfaces {
  ge-fpc/pic/port {
    accounting-profile name;
    description text;
    disable;
    ether-options {
      802.3ad {
        aex;
        (backup | primary);
        lacp {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      ieee-802-3az-eee;
      link-mode mode;
      (loopback | no-loopback);
      speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    optics-options {
      alarm alarm-type;
      warning alarm-type;
```



```
  wavelength nanometers;
}
traceoptions {
  flag flag;
}
(traps | no-traps);
unit logical-unit-number {
  accounting-profile name;
  arp-resp;
  bandwidth rate;
  description text;
  disable;
  family ccc;
  family ethernet-switching {
    filter {
      input filter-name;
      output filter-name;
    }
    native-vlan-id vlan-id-number;
    port-mode (access | trunk);
    vlan {
      members [ members];
    }
  }
  family inet {
    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;
        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 ];
        virtual-link-local-address address;
        vrrp-inherit-from {
          active-group group-number;
          active-interface interface-name;
        }
      }
    }
```

```
    dhcp {
      client-identifier (ascii client-id | hexadecimal client-id);
      lease-time (seconds | infinte);
      retransmission-attempt  number;
      retransmission-interval sections;
      server-address ip-address;
      update-server
      vendor-id
    }
    filter {
      input filter-name;
      output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    no-redirects;
    primary;
    rpf-check;
    targeted-broadcast;
  }
  family inet6 {
    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;
        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 ( address | routing-instance routing-instance-name );
        }
        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 group-name;
      input filter-name;
```

```
          output filter-name;
        }
        mtu bytes;
        no-neighbor-learn;
        policer {
          input policer-name;
          output policer-name;
        }
        rpf-check;
      }
      family iso {
        address interface-address;
        mtu bytes;
      }
      family mpls {
        mtu bytes;
      }
      proxy-arp (restricted | unrestricted);
      swap-by-poppush;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
}
```

Unsupported Statements in the [edit interfaces ge] Hierarchy Level

All statements in the **[edit interfaces ge]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153

[edit interfaces interface-range] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces interface-range]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

- Supported Statements in the [edit interfaces interface-range] Hierarchy Level on page 162
- Unsupported Statements in the [edit interfaces interface-range] Hierarchy Level on page 166

Supported Statements in the [edit interfaces interface-range] Hierarchy Level

The following hierarchy shows the **[edit interfaces interface-range]** configuration statements supported on EX Series switches.

```
interfaces {
interface-range name {
  accounting-profile name;
  aggregated-ether-options {
    ethernet-switch-profile {
      tag-protocol-id identifier ;
    }
    (flow-control | no-flow-control);
    ieee-802-3az-eee;
    lacp {
      (active | passive);
      admin-key key;
      periodic interval;
      system-id mac-address;
    }
    (link-protection | no-link-protection);
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
    rebalance-periodic;
    source-address-filter filter;
    source-filtering | no-source-filtering;
  }
  description text;
  disable;
  ether-options {
    802.3ad {
      aex;
      (backup | primary);
      lacp {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    link-mode mode;
    (loopback | no-loopback);
    speed (auto-negotiation | speed);
  }
  framing;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  member interface-name;
```



```
member-range starting-interface name to ending-interface name;
mtu bytes;
no-gratuitous-arp-request;
optics-options {
  alarm alarm-type;
  warning alarm-type;
  wavelength nanometers;
}
services-options;
speed speed;
traceoptions {
  flag flag;
}
(traps | no-traps);
unit logical-unit-number {
  accept-source-mac;
  accounting-profile name;
  arp-resp;
  bandwidth rate;
  description text;
  disable;
  family ccc;
  family ethernet-switching {
    filter {
      input filter-name;
      output filter-name;
    }
    native-vlan-id vlan-id-number;
    port-mode (access | trunk);
    vlan {
      members [ members];
    }
  }
  family inet {
    accounting {
      destination-class-usage;
      source-class-usage;
    }
    address ipv4-address {
      arp ip-address (mac | multicast-mac) mac-address <publish>;
      broadcast address;
      destination-class-usage;
      destination-profile;
      master-only;
      preferred;
      primary;
      vrrp-group group-number {
        (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 bandwidth;
              priority-cost number;
            }
            priority-hold-time seconds;
            route ip-address/mask routing-instance instance-name priority-cost cost;
          }
          virtual-address [ addresses ];
          virtual-link-local-address address;
          vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
          }
        }
      }
      dhcp {
        client-identifier (ascii client-id | hexadecimal client-id);
        lease-time (seconds | infinte);
        retransmission-attempt  number;
        retransmission-interval sections;
        server-address ip-address;
        update-server
        vendor-id
      }
      filter {
        input filter-name;
        output filter-name;
      }
      ipsec-sa;
      mtu bytes;
      multicast-only;
      negotiate-address;
      next-hop-tunnel;
      no-neighbor-learn;
      no-redirects;
      primary;
      receive-option-packets;
      rpf-check;
      targeted-broadcast;
    }
    family inet6 {
      accounting {
        destination-class-usage;
        source-class-usage;
      }
      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;
          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 ( address | routing-instance routing-instance-name );
                }
                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 group-name;
              input filter-name;
              output filter-name;
            }
            mtu bytes;
            no-neighbor-learn;
            policer {
              input policer-name;
              output policer-name;
            }
            rpf-check;
          }
          family iso {
            address interface-address;
            mtu bytes;
          }
          family mpls {
            mtu bytes;
          }
          interface-shared-with;
          interleave-fragments;
          inverse-arp;
          link-layer-overhead;
          minimum-links;
          mtu;
          proxy-arp (restricted | unrestricted);
          swap-by-poppush;
          (traps | no-traps);
          vlan-id vlan-id-number;
        }
        vlan-tagging;
      }
```

Unsupported Statements in the [edit interfaces interface-range] Hierarchy Level

All statements in the **[edit interfaces interface-range]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 34: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches

Statement	Hierarchy
NOTE: Variables, such as *interface-range*, are not shown in the statements or hierarchies.	
aggregated-sonet-options and all substatements	**[edit interfaces interface-range]**
allow-any-vci	**[edit interfaces interface-range unit]**
atm-l2circuit-mode	**[edit interfaces interface-range unit]**
atm-options and all substatements	**[edit interfaces interface-range]**
cell-bundle-size	**[edit interfaces interface-range unit]**
clear-don't-fragment-bit	**[edit interfaces interface-range unit]**
clocking	**[edit interfaces interface-range]**
compression-device	**[edit interfaces interface-range unit]**
container-options and all substatements	**[edit interfaces interface-range]**
copy-tos-to-outer-ip-header	**[edit interfaces interface-range unit]**
dce	**[edit interfaces interface-range]**
disable-mlppp-inter-ppp-pfc	**[edit interfaces interface-range unit]**
dlci	**[edit interfaces interface-range unit]**
drop-timeout	**[edit interfaces interface-range unit]**
ds0-options and all substatements	**[edit interfaces interface-range]**
e1-options and all substatements	**[edit interfaces interface-range]**
e3-options and all substatements	**[edit interfaces interface-range]**
epd-threshold	**[edit interfaces interface-range unit]**
family mlfr-end-to-end and all substatements	**[edit interfaces interface-range unit]**

Table 34: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches *(continued)*

Statement	Hierarchy
family mlfr-uni-uni and all substatements	**[edit interfaces interface-range unit]**
family mlppp and all substatements	**[edit interfaces interface-range unit]**
fragment-threshold	**[edit interfaces interface-range unit]**
ggsn-options and all substatements	**[edit interfaces interface-range]**
keepalives | no-keepalives	**[edit interfaces interface-range]** **[edit interfaces interface-range unit]**
lmi	**[edit interfaces interface-range]**
lsq-failure-options	**[edit interfaces interface-range]**
mlfr-uni-nni-bundle-options and all substatements	**[edit interfaces interface-range]**
mrru	**[edit interfaces interface-range unit]**
multicast-dlci	**[edit interfaces interface-range unit]**
multilink-max-classes	**[edit interfaces interface-range unit]**
multipoint	**[edit interfaces interface-range unit]**
multipoint-destination	**[edit interfaces interface-range unit family inet address]**
multiservice-options and all substatements	**[edit interfaces interface-range]**
oam-liveness	**[edit interfaces interface-range unit]**
oam-period	**[edit interfaces interface-range unit]**
passive-monitor-mode	**[edit interfaces interface-range unit]**
peer-unit	**[edit interfaces interface-range unit]**
plp-to-clp	**[edit interfaces interface-range unit]**
point-to-point	**[edit interfaces interface-range unit]**
ppp-options and all substatements	**[edit interfaces interface-range]** **[edit interfaces interface-range unit]**
receive-lsp	**[edit interfaces interface-range unit]**

Table 34: Unsupported [edit interfaces interface-range] Configuration Statements for EX Series Switches *(continued)*

Statement	Hierarchy
satop-options and all substatements	**[edit interfaces interface-range]**
serial-options and all substatements	**[edit interfaces interface-range]**
service-domain	**[edit interfaces interface-range unit]**
shaping	**[edit interfaces interface-range unit]**
short-sequence	**[edit interfaces interface-range unit]**
shdsl-options and all substatements	**[edit interfaces interface-range]**
t1-options and all substatements	**[edit interfaces interface-range]**
t3-options and all substatements	**[edit interfaces interface-range]**
transmit-lsp	**[edit interfaces interface-range unit]**
transmit-weight	**[edit interfaces interface-range unit]**
trunk-id	**[edit interfaces interface-range unit]**
tunnel	**[edit interfaces interface-range unit]**
vci	**[edit interfaces interface-range unit]**
vci-range	**[edit interfaces interface-range unit]**
vpi	**[edit interfaces interface-range unit]**

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153

[edit interfaces lo] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces lo]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

- Supported Statements in the [edit interfaces lo] Hierarchy Level on page 169
- Unsupported Statements in the [edit interfaces lo] Hierarchy Level on page 171

Supported Statements in the [edit interfaces lo] Hierarchy Level

The following hierarchy shows the **[edit interfaces lo]** configuration statements supported on EX Series switches.

```
interfaces {
  lo0 {
    accounting-profile name;
    description text;
    disable;
    hold-time down milliseconds up milliseconds ;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ccc;
      family inet {
        address ipv4-address {
          preferred;
          primary;
          vrrp-group group-number {
            (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 bandwidth;
                priority-cost number;
              }
              priority-hold-time seconds;
              route ip-address/mask routing-instance instance-name priority-cost cost;
            }
            virtual-address [ addresses ];
            virtual-link-local-address address;
            vrrp-inherit-from {
              active-group group-number;
              active-interface interface-name;
            }
```

```
      }
    }
    dhcp {
      client-identifier (ascii client-id | hexadecimal client-id);
      lease-time (seconds | infinte);
      retransmission-attempt  number;
      retransmission-interval sections;
      server-address ip-address;
      update-server
      vendor-id
    }
    filter {
      input filter-name;
      output filter-name;
    }
    no-neighbor-learn;
    no-redirects;
    primary;
  }
  family inet6 {
    address address {
      preferred;
      primary;
      vrrp-inet6-group group-id {
        accept-data | no-accept-data;
        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 ( address | routing-instance routing-instance-name );
        }
        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 group-name;
      input filter-name;
      output filter-name;
    }
    no-neighbor-learn;
    policer {
```

```
            input policer-name;
            output policer-name;
          }
        }
        family iso {
          address interface-address;
        }
        family mpls;
        (traps | no-traps);
      }
    }
  }
```

Unsupported Statements in the [edit interfaces lo] Hierarchy Level

All statements in the **[edit interfaces lo]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 35: Unsupported [edit interfaces lo] Configuration Statements for EX Series Switches

Statement	Hierarchy
layer2-policer	[edit interfaces lo unit]
any	[edit interfaces lo unit family]
tcc	[edit interfaces lo unit family]
policer	[edit interfaces lo unit family inet]
unnumbered-address	[edit interfaces lo unit family inet]

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153
- *[edit interfaces] Configuration Statement Hierarchy on EX Series Switches*

[edit interfaces me] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces me]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

- Supported Statements in the [edit interfaces me] Hierarchy Level on page 172
- Unsupported Statements in the [edit interfaces me] Hierarchy Level on page 174

Supported Statements in the [edit interfaces me] Hierarchy Level

The following hierarchy shows the **[edit interfaces me]** configuration statements supported on EX Series switches.

```
interfaces {
  me0 {
    accounting-profile name;
    description text;
    disable;
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions {
      flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
      accounting-profile name;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      family ethernet-switching {
        filter {
          input filter-name;
          output filter-name;
        }
        native-vlan-id vlan-id-number;
        port-mode (access | trunk);
        vlan {
          members [ members];
        }
      }
      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;
          master-only;
          preferred;
          primary;
        }
        dhcp {
```

```
        client-identifier (ascii client-id | hexadecimal client-id);
        lease-time (seconds | infinte);
        retransmission-attempt  number;
        retransmission-interval sections;
        server-address ip-address;
        update-server
        vendor-id
      }
      filter {
        input filter-name;
        output filter-name;
      }
      mtu bytes;
      no-neighbor-learn;
      primary;
      rpf-check;
    }
    family inet6 {
      accounting {
        destination-class-usage;
        source-class-usage {
          input;
          output;
        }
      }
      address address {
        eui-64;
        ndp ip-address (mac | multicast-mac) mac-address &lt;publish&gt;;
        preferred;
        primary;
      }
      (dad-disable | no-dad-disable);
      filter {
        group group-name;
        input filter-name;
        output filter-name;
      }
      mtu bytes;
      no-neighbor-learn;
      policer {
        input policer-name;
        output policer-name;
      }
      rpf-check;
    }
    family iso {
      address interface-address;
      mtu bytes;
    }
    family mpls {
      mtu bytes;
    }
    swap-by-poppush;
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
```

```
        vlan-tagging;
    }
}
```

Unsupported Statements in the [edit interfaces me] Hierarchy Level

All statements in the **[edit interfaces me]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 36: Unsupported [edit interfaces me] Configuration Statements for EX Series Switches

Statement	Hierarchy
encapsulation	[edit interfaces me]
link-mode	[edit interfaces me]
encapsulation	[edit interfaces me unit]
layer2-policer	[edit interfaces me unit]
native-inner-vlan-id	[edit interfaces me unit]
vlan-id-list	[edit interfaces me unit]
vlan-id-range	[edit interfaces me unit]
ccc	[edit interfaces me unit family]
tcc	[edit interfaces me unit family]
vpls	[edit interfaces me unit family]
no-redirects	[edit interfaces me unit family inet]
policer	[edit interfaces me unit family inet]
sampling	[edit interfaces me unit family inet]
service	[edit interfaces me unit family inet]
unnumbered-address	[edit interfaces me unit family inet]
vrrp-group	[edit interfaces me unit family inet address]
service	[edit interfaces me unit family inet6]
vrrp-inet6-group	[edit interfaces me unit family inet6 address]

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153

- *[edit interfaces] Configuration Statement Hierarchy on EX Series Switches*

[edit interfaces vlan] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces vlan]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

- Supported Statements in the [edit interfaces vlan] Hierarchy Level on page 175
- Unsupported Statements in the [edit interfaces vlan] Hierarchy Level on page 178

Supported Statements in the [edit interfaces vlan] Hierarchy Level

The following hierarchy shows the **[edit interfaces vlan]** configuration statements supported on EX Series switches.

```
interfaces {
  vlan {
    accounting-profile name;
    description  text;
    disable;
    (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;
      arp-resp;
      bandwidth rate;
      description text;
      disable;
      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;
      master-only;
      preferred;
      primary;
      vrrp-group group-number {
        (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 bandwidth;
            priority-cost number;
          }
          priority-hold-time seconds;
          route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
        virtual-link-local-address address;
        vrrp-inherit-from {
          active-group group-number;
          active-interface interface-name;
        }
      }
    }
    dhcp {
      client-identifier (ascii client-id | hexadecimal client-id);
      lease-time (seconds | infinte);
      retransmission-attempt  number;
      retransmission-interval sections;
      server-address ip-address;
      update-server
      vendor-id
    }
    filter {
      input filter-name;
      output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    primary;
    rpf-check;
  }
  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;
        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 ( address | routing-instance routing-instance-name );
        }
        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 group-name;
      input filter-name;
      output filter-name;
    }
    mtu bytes;
    no-neighbor-learn;
    policer {
      input policer-name;
      output policer-name;
    }
    rpf-check;
  }
  family iso {
    address interface-address;
    mtu bytes;
  }
  family mpls {
    mtu bytes;
  }
  proxy-arp (restricted | unrestricted);
```

```
            (traps | no-traps);
        }
    }
}
```

Unsupported Statements in the [edit interfaces vlan] Hierarchy Level

All statements in the **[edit interfaces vlan]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153

[edit interfaces vme] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces vme]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

Supported Statements in the [edit interfaces vme] Hierarchy Level

The following hierarchy shows the **[edit interfaces vme]** configuration statements supported on EX Series switches.

```
interfaces {
    vme {
        accounting-profile name;
        description  text;
        disable;
        (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;
            arp-resp;
            bandwidth rate;
```



```
description text;
disable;
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;
    master-only;
    preferred;
    primary;
    vrrp-group group-number {
      (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 bandwidth;
          priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
      }
      virtual-address [ addresses ];
      virtual-link-local-address address;
      vrrp-inherit-from {
        active-group group-number;
        active-interface interface-name;
      }
    }
  }
  dhcp {
    client-identifier (ascii client-id | hexadecimal client-id);
    lease-time (seconds | infinte);
    retransmission-attempt  number;
    retransmission-interval sections;
    server-address ip-address;
    update-server
    vendor-id
  }
  filter {
    input filter-name;
    output filter-name;
  }
  mtu bytes;
```

```
        no-neighbor-learn;
        primary;
        rpf-check;
      }
      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;
            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 ( address | routing-instance routing-instance-name );
            }
            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 group-name;
          input filter-name;
          output filter-name;
        }
        mtu bytes;
        no-neighbor-learn;
        policer {
          input policer-name;
          output policer-name;
        }
        rpf-check;
```



```
            }
            family iso {
                address interface-address;
                mtu bytes;
            }
            family mpls {
                mtu bytes;
            }
            (traps | no-traps);
            vlan-id vlan-id-number;
        }
        vlan-tagging;
    }
}
```

Unsupported Statements in the [edit interfaces vme] Hierarchy Level

All statements in the **[edit interfaces vme]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153

- *[edit interfaces] Configuration Statement Hierarchy on EX Series Switches*

[edit interfaces xe] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit interfaces xe]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see Feature Explorer.

This topic lists:

- Supported Statements in the [edit interfaces xe] Hierarchy Level on page 181
- Unsupported Statements in the [edit interfaces xe] Hierarchy Level on page 185

Supported Statements in the [edit interfaces xe] Hierarchy Level

The following hierarchy shows the **[edit interfaces xe]** configuration statements supported on EX Series switches.

```
interfaces {
    xe-fpc/pic/port {
        accounting-profile name;
        clocking (external | internal);
        description text;
```

```
disable;
ether-options {
   802.3ad {
      aex;
      (backup | primary);
      lacp {
         force-up;
      }
   }
   (flow-control | no-flow-control);
   (loopback | no-loopback);
}
framing (lan-phy | wan-phy);
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
mtu bytes;
no-gratuitous-arp-request;
optics-options {
   alarm alarm-type;
   warning alarm-type;
   wavelength nanometers;
}
traceoptions {
   flag flag;
}
(traps | no-traps);
unit logical-unit-number {
   accounting-profile name;
   bandwidth rate;
   description text;
   disable;
   family ccc;
   family ethernet-switching {
      filter {
         input filter-name;
         output filter-name;
      }
      native-vlan-id vlan-id-number;
      port-mode (access | trunk);
      vlan {
         members [ members];
      }
   }
   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;
          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 ];
          virtual-link-local-address address;
          vrrp-inherit-from {
            active-group group-number;
            active-interface interface-name;
          }
        }
      }
      dhcp {
        client-identifier (ascii client-id | hexadecimal client-id);
        lease-time (seconds | infinte);
        retransmission-attempt  number;
        retransmission-interval sections;
        server-address ip-address;
        update-server
        vendor-id
      }
      filter {
        input filter-name;
        output filter-name;
      }
      mtu bytes;
      no-neighbor-learn;
      no-redirects;
      primary;
      rpf-check;
      targeted-broadcast;
    }
    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;
                    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 ( address | routing-instance routing-instance-name );
                    }
                    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 group-name;
                input filter-name;
                output filter-name;
            }
            mtu bytes;
            no-neighbor-learn;
            policer {
                input policer-name;
                output policer-name;
            }
            rpf-check;
        }
        family iso {
            address interface-address;
            mtu bytes;
        }
        family mpls {
            mtu bytes;
        }
        proxy-arp (restricted | unrestricted);
        swap-by-poppush;
        (traps | no-traps);
        vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
    }
    vlan-tagging;
```

```
  }
}
```

Unsupported Statements in the [edit interfaces xe] Hierarchy Level

All statements in the **[edit interfaces xe]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- [edit interfaces] Configuration Statement Hierarchy on EX Series Switches on page 153

[edit protocols lacp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit protocols lacp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*.

This topic lists:

- Supported Statements in the [edit protocols lacp] Hierarchy Level on page 185
- Unsupported Statements in the [edit protocols lacp] Hierarchy Level on page 186

Supported Statements in the [edit protocols lacp] Hierarchy Level

The following hierarchy shows the **[edit protocols lacp]** configuration statements supported on EX Series switches:

```
protocols {
  lacp {
    ppm {
      centralized
    }
    traceoptions {
      file <filename> <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
}
```

Unsupported Statements in the [edit protocols lacp] Hierarchy Level

All statements in the **[edit protocols lacp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- *[edit protocols] Configuration Statement Hierarchy on EX Series Switches*

802.3ad

Syntax

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

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

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure membership in a link aggregation group (LAG).

Options

- **ae*x***—Name of the LAG.
- **backup**—Designate the interface as the backup interface for link-protection mode.
- **primary**—Designate the interface as the primary interface for link-protection mode.

The remaining statements are described separately.

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

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches on page 59
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

accounting-profile

Syntax accounting-profile *name*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Enable collection of accounting data for the specified physical or logical interface or interface range.

Options ***name***—Name of the accounting profile.

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

Related Documentation

- Applying an Accounting Profile to the Physical Interface on page 113
- Applying an Accounting Profile to the Logical Interface on page 114

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 dlci dlci-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 84.

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 on page 208
- *negotiate-address*
- *unnumbered-address (Ethernet)*

aggregated-devices

Syntax

```
aggregated-devices {
  ethernet (Aggregated Devices) {
    device-count number;
    lacp
  }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure properties for aggregated devices on the switch.

The remaining statements are explained separately.

Default Aggregated devices are disabled.

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

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- Understanding Aggregated Ethernet Interfaces and LACP on page 8
- *Junos OS Ethernet Interfaces Configuration Guide*

aggregated-ether-options

Syntax

```
aggregated-ether-options {
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    admin-key key;
    periodic interval;
    system-id mac-address;
  }
  (link-protection | no-link-protection);
  link-protection-sub-group group-name {
    [primary | backup);
  }
  link-speed speed;
  (loopback | no-loopback);
  minimum-links number;
  system-priority
}
```

Hierarchy Level [edit interfaces ae*x*]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.

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

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- Understanding Aggregated Ethernet Interfaces and LACP on page 8
- *Junos OS Ethernet Interfaces Configuration Guide*

alarm (optics-options)

Syntax

```
alarm low-light–alarm {
  (link-down | syslog);
}
```

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Specify the action to take if the receiving optics signal is below the optics low-light alarm threshold.

Options **link-down**—Drop the 10-Gigabit Ethernet link and marks link as down.

syslog—Write the optics information to the system log.

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

Related Documentation

- *Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning*
- *100-Gigabit Ethernet OTN Options Configuration Overview*

arp (Interfaces)

Syntax

arp *ip-address* (mac | multicast-mac) *mac-address* publish;

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.

Options

ip-address—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing **address** statement.

mac *mac-address*—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**.

multicast-mac *mac-address*—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***. For example, **0011.2233.4455** or **00:11:22:33:44:55**.

publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.

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 Static ARP Table Entries on page 117

auto-negotiation

Syntax

(auto-negotiation | no-auto-negotiation) <remote-fault (local-interface-online | local-interface-offline)>;

Hierarchy Level

[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information

Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 8.4 for J Series Services Routers.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description

For Gigabit Ethernet interfaces on M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers explicitly enable autonegotiation and remote fault. For EX Series switches and J Series Services Routers, explicitly enable autonegotiation only.

- **auto-negotiation**—Enables autonegotiation. This is the default.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure the link mode and speed.

When you configure Tri-Rate Ethernet copper interfaces to operate at 1 Gbps, autonegotiation must be enabled.

NOTE: On EX Series switches, an interface configuration that disables autonegotiation and manually sets the link speed to 1 Gbps is accepted when you commit the configuration; however, if the interface you are configuring is a Tri-Rate Ethernet copper interface, the configuration is ignored as invalid and autonegotiation is enabled by default.

To correct the invalid configuration and disable autonegotiation:

1. Delete the **no-auto-negotiation** statement and commit the configuration.
2. Set the link speed to 10 or 100 Mbps, set **no-auto-negotiation**, and commit the configuration.

On J Series Services Routers with universal Physical Interface Modules (uPIMs) and on EX Series switches, if the link speed and duplex mode are also configured, the interfaces use the values configured as the desired values in the negotiation. If autonegotiation is disabled, the link speed and link mode must be configured.

NOTE: On T4000 routers, the **auto-negotiation** command is ignored for interfaces other than Gigabit Ethernet.

Default Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Options **remote-fault (local-interface-online | local-interface-offline)**—(Optional) For M Series, MX Series, T Series, TX Matrix routers, and ACX Series routers only, manually configure remote fault on an interface.

Default: local-interface-online

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

Related Documentation

- *Gigabit Ethernet Autonegotiation Overview*
- *Configuring Gigabit Ethernet Interfaces on J Series Services Routers*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

bandwidth (Interfaces)

Syntax bandwidth *rate*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure an informational-only bandwidth value for an interface. This statement is valid for all logical interface types except multilink and aggregated interfaces.

NOTE: We recommend that you be careful when setting this value. Any interface bandwidth value that you configure using the **bandwidth** statement affects how the interface cost is calculated for a dynamic routing protocol, such as OSPF. By default, the interface cost for a dynamic routing protocol is calculated using the following formula:

cost = reference-bandwidth/bandwidth,

where bandwidth is the physical interface speed. However, if you specify a value for bandwidth using the **bandwidth** statement, that value is used to calculate the interface cost, rather than the actual physical interface bandwidth.

Options ***rate***—Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation **c**; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.

Range: Not limited.

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

Related Documentation

- Configuring the Interface Bandwidth on page 89

broadcast

Syntax broadcast *address*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0, nor can you specify a broadcast address.

Default The default broadcast address has a host portion of all ones.

Options *address*—Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses **0.0.0.0** or **255.255.255.255**.

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 Interface Address on page 84

chassis

Syntax

```
chassis {
  aggregated-devices {
    ethernet (Aggregated Devices) {
      device-count number;
    }
  }
  auto-image-upgrade;
  fpc slot {
    pic pic-number {
      sfpplus {
        pic-mode mode;
      }
    }
    power-budget-priority priority;
  }
  lcd-menu {
    fpc slot-number {
      menu-item (menu-name | menu-option) {
        disable;
      }
    }
  }
  nssu {
    upgrade-group group-name {
      fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
      member (NSSU Upgrade Groups) member-id {
        fpcs (NSSU Upgrade Groups) (slot-number | [list-of-slot-numbers]);
      }
    }
  }
  psu {
    redundancy {
      n-plus-n (Power Management);
    }
  }
  redundancy {
    graceful-switchover;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

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 Aggregated Ethernet Links (CLI Procedure) on page 121
- *Upgrading Software by Using Automatic Software Download*
- *Configuring the LCD Panel on EX Series Switches (CLI Procedure)*
- *Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)*
- *Configuring Power Supply Redundancy (CLI Procedure)*
- *Configuring the Power Priority of Line Cards (CLI Procedure)*
- *Configuring Line-Card Upgrade Groups for Nonstop Software Upgrade (CLI Procedure)*

description (Interfaces)

Syntax description *text*;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the **show interfaces** commands, and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no effect on the operation of the interface on the router or switch.

The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.

Options *text*—Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.

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

Related Documentation

- *Configuring Interface Description*
- Adding a Logical Unit Description to the Configuration on page 81
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- *Using DHCP Relay Agent Option 82 Information*
- *Junos OS Network Interfaces Library for Routing Devices*
- *Example: Connecting Access Switches to a Distribution Switch*

destination (Tunnels)

Syntax destination *address*;

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

Release Information 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.

Description For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.

Options ***address***—Address of the remote side of the connection.

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

Related Documentation

- Configuring the Interface Address on page 84
- *point-to-point*

device-count

Syntax device-count *number*;

Hierarchy Level [edit chassis aggregated-devices ethernet (Aggregated Devices)]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Range updated in Junos OS Release 9.5 for EX Series switches.

Description Configure the number of aggregated Ethernet logical devices available to the switch.

Options *number*—Maximum number of aggregated Ethernet logical interfaces on the switch.

Range: 1 through 32 for EX2200, EX3200, and standalone EX3300 switches and for EX3300 Virtual Chassis

Range: 1 through 64 for standalone EX4200, standalone EX4500, and EX6200 switches and for EX4200 and EX4500 Virtual Chassis

Range: 1 through 239 for EX8200 Virtual Chassis

Range: 1 through 255 for standalone EX8200 switches

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

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- *Junos OS Network Interfaces Configuration Guide*

disable (Interface)

Syntax disable;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description Disable a physical or a logical interface, effectively unconfiguring it.

CAUTION:

- Dynamic subscribers and logical interfaces use physical interfaces for connection to the network. The Junos OS allows you to set the interface to disable and commit the change while dynamic subscribers and logical interfaces are still active. This action results in the loss of all subscriber connections on the interface. Use care when disabling interfaces.
- If aggregated SONET links are configured between a T1600 router and a T4000 router, interface traffic is disrupted when you disable the physical interface configured on the T1600 router. If you want to remove the interface, we recommend that you deactivate the interface instead of disabling it.

NOTE:

- When you use the **disable** statement at the **[edit interfaces]** hierarchy level, depending on the PIC type, the interface might or might not turn off the laser. Older PIC transceivers do not support turning off the laser, but newer Gigabit Ethernet (GE) PICs with SFP and XFP transceivers and ATM MIC with SFP do support it and the laser will be turned off when the interface is disabled. If the ATM MIC with SFP is part of an APS group, then the laser will not be turned off when you use the **disable** statement at the **[edit interfaces]** hierarchy level..
- When you disable or deactivate an interface, then all the references made to the deactivated interface must be removed from the routing instance.

WARNING: Do not stare into the laser beam or view it directly with optical instruments even if the interface has been disabled.

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

Related Documentation

- Disabling a Physical Interface on page 82
- Disabling a Logical Interface on page 83

disable (Link Protection)

Syntax disable;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Disable LACP link protection on the interface.

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

Related Documentation

- *Configuring LACP for Aggregated Ethernet Interfaces*
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

disable (Multicast Load Balancing)

Syntax disable;

Hierarchy Level [edit chassis multicast-loadbalance]

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

Description (EX8200 switches only) Disable multicast load balancing. After you configure this statement and commit it, multicast load balancing no longer balances traffic on aggregated 10-gigabit Ethernet interfaces that are currently configured as well as those that are configured later.

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

Related Documentation

- Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure) on page 144

eui-64

Syntax eui-64;

Hierarchy Level [edit interfaces *interface-name* unit *number* family inet6 address *address*]

Release Information 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.

Description For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.

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

Related Documentation

- Configuring the Interface Address on page 84

ether-options

Syntax Gigabit Ethernet interfaces:

```
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
      port-priority
    }
  }
  (auto-negotiation | no-auto-negotiation);
  (flow-control | no-flow-control);
  ieee-802-3az-eee;
  ignore-l3-incompletes;
  link-mode mode;
  (loopback | no-loopback);
  speed (speed | auto-negotiation);
}
```

10-Gigabit Ethernet interfaces:

```
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
      port-priority
    }
  }
  (flow-control | no-flow-control);
  (loopback | no-loopback);
}
```

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

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure Ethernet properties for a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface on an EX Series 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 Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70

- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- Understanding Aggregated Ethernet Interfaces and LACP on page 8
- EX Series Switches Interfaces Overview on page 3
- *Junos OS Ethernet Interfaces Configuration Guide*

ethernet (Aggregated Devices)

Syntax

```
ethernet {
device-count number;
lacp {
  link-protection {
    non-revertive;
  }
  system-priority;
}
```

Hierarchy Level [edit chassis aggregated-devices]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure properties for Ethernet aggregated devices on the switch.

The remaining statement is 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 Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- *Junos OS Ethernet Interfaces Configuration Guide*

family (for EX Series switches)

Syntax

family ccc on page 208
family ethernet-switching on page 208
family inet on page 208
family inet6 on page 209
family iso on page 209
family mpls on page 209

family ccc

```
family ccc;
```

family ethernet-switching

```
family ethernet-switching {
  filter [input | output] filter-name;
  native-vlan-id vlan-id;
  port-mode mode;
  vlan (802.1Q Tagging) {
    members [ (all | names | vlan-ids)];
  }
}
```

family inet

```
family inet {
  address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast;
    preferred;
    primary;
    vrrp-group group-id {
      advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      virtual-address [addresses];
      virtual-link-local-address ip-address;
    }
  }
  dhcp {
    client-identifier (ascii ascii | hexadecimal hexadecimal);
    lease-time (seconds | infinite);
    retransmission-attempt  number;
    retransmission-interval seconds;
    server-address ip-address;
    update-server;
    vendor-id vendor-id;
  }
  filter {
    input filter-name;
    output filter-name;
    }
  mtu bytes;
  no-redirects;
  no-neighbor-learn;
  primary;
  rpf-check;
```

```
    targeted-broadcast;
}
```

family inet6

```
family inet6 {
    address address {
        eui-64;
        nd6-stale-time seconds;
        ndp ip-address (mac | multicast-mac) mac-address <publish>;
        preferred;
        primary;
        vrrp-inet6-group group-id {
            inet6-advertise-interval milliseconds;
            preempt | preempt {
                hold-time seconds;
            }
            priority number;
            virtual-inet6-address [addresses];
            virtual-link-local-address ipv6–address;
        }
    }
    (dad-disable | no-dad-disable);
    filter {
        input filter-name;
        output filter-name;
    }
    mtu bytes;
    no-neighbor-learn
    rpf-check;
}
```

family iso

```
family iso {
    address interface-address;
    mtu bytes;
}
```

family mpls

```
family mpls {
    mtu bytes;
}
```

Hierarchy Level

[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *name* unit *logical-unit-number*]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches, including options **ethernet-switching**, **inet**, and **iso**.
Option **inet6** introduced in Junos OS Release 9.3 for EX Series switches.
Options **ccc** and **mpls** introduced in Junos OS Release 9.5 for EX Series switches.

Description

Configure protocol family information for the logical interface on the switch.

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

Default Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. Before you can change the family setting for an interface to another family type, you must delete this default setting or any user-configured family setting. EX6200 and EX8200 switch interfaces do not have a default family setting.

Options See Table 37 on page 210 for protocol families available on the switch interfaces. Different protocol families support different subsets of the interface types on the switch. Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Interface-range configuration (**interface-range**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)
- Virtual management Ethernet (**vme**)
- 10-Gigabit Ethernet (**xe**)

If you are using an interface range, the supported protocol families are the ones supported by the interface types that compose the range.

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

Table 37: Protocol Families and Supported Interface Types

Family	Description	Supported Interface Types						
		ae	ge	lo0	me0	vlan	vme	xe
ccc	Circuit cross-connect protocol family	✓*	✓					✓
ethernet-switching	Ethernet switching protocol family	✓	✓		✓			✓
inet	IPv4 protocol family	✓	✓	✓	✓	✓	✓	✓
inet6	IPv6 protocol family	✓	✓	✓	✓	✓	✓	✓
iso	Junos OS protocol family for IS-IS traffic	✓	✓	✓	✓	✓	✓	✓
mpls	MPLS protocol family	✓	✓	✓	✓		✓	✓

*Supported on EX8200 switches only

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 a DHCP Server on Switches (CLI Procedure)*
- *Example: Configuring MPLS on EX8200 and EX4500 Switches*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- *Configuring Routed VLAN Interfaces (CLI Procedure)*

filter

Syntax

```
filter {
  group filter-group-number;
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
```

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 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description

Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the **family ethernet-switching**, **inet**, **inet6**, **mpls**, or **vpls** only.

NOTE: On QFX3500 and QFX3600 switches running Enhanced Layer 2 Software and on OCX Series switches, VPLS is not supported.

Options

group *filter-group-number*—Define an interface to be part of a filter group.

Range: 1 through 255

input *filter-name*—Name of one filter to evaluate when packets are received on the interface.

output *filter-name*—Name of one filter to evaluate when packets are transmitted on the interface.

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

- *Applying a Filter to an Interface*
- *Junos OS Services Interfaces Library for Routing Devices*
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*
- *Junos OS Administration Library for Routing Devices*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- *Configuring Firewall Filters (CLI Procedure)*
- *Configuring Firewall Filters and Policers for VPLS*
- family on page 208
- *family*

flow-control

Syntax (flow-control | no-flow-control);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces *interface-name* multiservice-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router or switch to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.

NOTE: On the Type 5 FPC, to prioritize control packets in case of ingress oversubscription, you must ensure that the neighboring peers support MAC flow control. If the peers do not support MAC flow control, then you must disable flow control.

Default Flow control is enabled.

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

Related Documentation

- Configuring Flow Control on page 84
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

force-up

Syntax force-up;

Hierarchy Level [edit interfaces *interface-name* ether-options 802.3ad lacp]

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.

Description Set the state of the interface as UP when the peer has limited LACP capability.

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

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- Understanding Aggregated Ethernet Interfaces and LACP on page 8
- *Junos OS Ethernet Interfaces Configuration Guide*

gratuitous-arp-reply

Syntax (gratuitous-arp-reply | no-gratuitous-arp-reply);

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 in EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For Ethernet interfaces, enable updating of the Address Resolution Protocol (ARP) cache for gratuitous ARPs.

Default Updating of the ARP cache is disabled on all Ethernet interfaces.

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

Related Documentation

- Configuring Gratuitous ARP on page 116
- *no-gratuitous-arp-request*

hold-time (Physical Interface)

Syntax hold-time up *milliseconds* down *milliseconds*;

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

Release Information 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.

Description Specify the **hold-time** 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.

NOTE:

- 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.
- The **hold-time** option is not available for controller interfaces.

Default Interface transitions are not damped.

Options **down** ***milliseconds***—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.

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

up ***milliseconds***—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.

Range: 0 through 4,294,967,295

Default: 0 (interface transitions are not damped)

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

Related Documentation

- *advertise-interval*
- interfaces (for EX Series switches) on page 219

- *Physical Interface Damping Overview*
- Damping Shorter Physical Interface Transitions on page 146
- *Damping Longer Physical Interface Transitions*

ieee-802-3az-eee

Syntax ieee-802-3az-eee;

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

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

Description Configure Energy Efficient Ethernet (EEE) on an EEE-capable Base-T copper interface.

Default EEE is disabled on EEE-capable interfaces.

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

Related Documentation

- Configuring Energy Efficient Ethernet on Interfaces (CLI Procedure) on page 145

interface-range

Syntax

```
interface-range name {
  accounting-profile name;
  description text;
  disable;
  ether-options {
    802.3ad {
      aex;
      (backup | primary);
      lacp {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    ieee-802-3az-eee;
    link-mode mode;
    (loopback | no-loopback);
    speed (auto-negotiation | speed);
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  member interface-name;
  member-range  starting-interface name to  ending-interface name;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
    vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
  }
  vlan-tagging;
}
```

Hierarchy Level [edit interfaces]

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.

Description Group interfaces that share a common configuration profile.

NOTE: You can specify interface ranges only for Gigabit and 10-Gigabit Ethernet interfaces.

Options *name*—Name of the interface range.

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

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 Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Understanding Interface Ranges on EX Series Switches on page 17
- *Understanding Interface Ranges on EX Series Switches*
- EX Series Switches Interfaces Overview on page 3
- *Junos OS Interfaces Fundamentals Configuration Guide*

interfaces (for EX Series switches)

Syntax

interfaces ae on page 219
interfaces ge on page 219
interfaces interface-range on page 221
interfaces lo0 on page 221
interfaces me0 on page 222
interfaces traceoptions on page 222
interfaces vlan on page 222
interfaces vme on page 223
interfaces xe on page 224

interfaces ae

```
aex {
  accounting-profile name;
  aggregated-ether-options {
    (flow-control | no-flow-control);
    lacp {
      (active | passive);
      admin-key key;
      periodic interval;
      system-id mac-address;
    }
    (link-protection | no-link-protection);
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
  }
  description text;
  disable;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

interfaces ge

```
ge-fpc/pic/port {
  accounting-profile name;
  description text;
  disable;
  ether-options {
    802.3ad {
```

```
            aex;
            (backup | primary);
            lacp {
                force-up;
            }
        }
        (auto-negotiation | no-auto-negotiation);
        (flow-control | no-flow-control);
        ieee-802-3az-eee;
        link-mode mode;
        (loopback | no-loopback);
        speed (auto-negotiation | speed);
    }
    (gratuitous-arp-reply | no-gratuitous-arp-reply);
    hold-time up milliseconds down milliseconds;
    media-type;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions {
        flag flag;
    }
    (traps | no-traps);
    unit logical-unit-number {
        accounting-profile name;
        bandwidth rate;
        description text;
        disable;
        family family-name {...}
        proxy-arp (restricted | unrestricted);
        (traps | no-traps);
        vlan-id vlan-id-number;
    }
    vlan-tagging;
}
```

interfaces interface-range

```
interface-range name {
  accounting-profile name;
  description text;
  disable;
  ether-options {
    802.3ad {
      aex;
      (backup | primary);
      lacp {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    ieee-802-3az-eee;
    link-mode mode;
    (loopback | no-loopback);
    speed (auto-negotiation | speed);
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  member interface-name;
  member-range starting-interface name to ending-interface name;
  mtu bytes;
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

interfaces lo0

```
lo0 {
  accounting-profile name;
  description text;
  disable;
  hold-time up milliseconds down milliseconds;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
  }
}
```

interfaces me0

```
me0 {
  accounting-profile name;
  description text;
  disable;
  (gratuitous-arp-reply | no-gratuitous-arp-reply);
  hold-time up milliseconds down milliseconds;
  no-gratuitous-arp-request;
  traceoptions {
    flag flag;
  }
  (traps | no-traps);
  unit logical-unit-number {
    accounting-profile name;
    bandwidth rate;
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

interfaces traceoptions

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}
```

interfaces vlan

```
vlan {
  accounting-profile name;
  description  text;
  disable;
  (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;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
  }
}
```

interfaces vme

```
vme {
  accounting-profile name;
  description text;
  disable;
  (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;
    family family-name {...}
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```


interfaces xe

```
xe-fpc/pic/port {
  accounting-profile name;
  description text;
  disable;
  ether-options {
    802.3ad {
      aex;
      (backup | primary);
      lacp {
        force-up;
      }
    }
    (flow-control | no-flow-control);
    link-mode mode;
    (loopback | no-loopback);
  }
  (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;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure interfaces on EX Series switches.

Options See Table 38 on page 225 for the interface types and protocol-family options supported on the switch. Different protocol families support different subsets of the interface types on the switch. See the **family** statement for syntax of the protocol families supported for switch interfaces.

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

Table 38: Interface Types and Their Supported Protocol Families

		Supported Protocol Families					
Interface Type	Description	ccc	ethernet-switching	inet	inet6	iso	mpls
ae	Aggregated Ethernet interface (also referred to as a link aggregation group [LAG])	✓*	✓	✓	✓	✓	✓
ge	Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
interface-range	Interface-range configuration	Supported protocol families are the ones supported by the interface types that compose the range.					
lo0	Loopback interface			✓	✓	✓	✓
me0	Management Ethernet interface		✓	✓	✓	✓	✓
vlan	Routed VLAN interface (RVI)			✓	✓	✓	
vme	Virtual management Ethernet interface			✓	✓	✓	✓
xe	10-Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
		*Supported on EX8200 switches only					

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 Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 135
- *Configuring Routed VLAN Interfaces (CLI Procedure)*
- *Configuring the Virtual Management Ethernet Interface for Global Management of an EX Series Virtual Chassis (CLI Procedure)*
- EX Series Switches Interfaces Overview on page 3
- *Junos OS Interfaces Fundamentals Configuration Guide*
- *Junos OS Ethernet Interfaces Configuration Guide*

lacp (Aggregated Ethernet)

Syntax

```
lacp {
  (active | passive);
  admin-key key;
  accept-data;
  fast-failover;
  link-protection {
    disable;
    (revertive |non-revertive);
  }
  periodic interval;
  system-id mac-address;
  system-priority priority;
}
```

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]
[edit logical-systems *logical-system-name* interfaces ae*X* aggregated-ether-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
fast-failover option introduced in Junos OS Release 12.2.
Support for logical systems introduced in Junos OS Release 14.1.

Description

Configure the Link Aggregation Control Protocol (LACP) for aggregated Ethernet interfaces only.

When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, the router processes packets received on a member link irrespective of the LACP state if the aggregated Ethernet bundle is up.

NOTE: When you configure the **accept-data** statement at the **[edit interfaces ae*X* aggregated-ether-options lacp]** hierarchy level, this behavior occurs:

- By default, the **accept-data** statement is not configured when LACP is enabled.
- You can configure the **accept-data** statement to improve convergence and reduce the number of dropped packets when member links in the bundle are enabled or disabled.
- When LACP is down and a member link receives packets, the router or switch does not process packets as defined in the IEEE 802.1ax standard. According to this standard, the packets should be dropped, but they are processed instead because the **accept-data** statement is configured.

Default

If you do not specify LACP as either **active** or **passive**, LACP remains passive.

Options

active—Initiate transmission of LACP packets.

admin-key ***number***—Specify an administrative key for the router or switch.

NOTE: You must also configure multichassis link aggregation (MC-LAG) when you configure the **admin-key**.

fast-failover—Specify to override the IEEE 802.3ad standard and allow the standby link to receive traffic. Overriding the default behavior facilitates subsecond failover.

passive—Respond to LACP packets.

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 LACP for Aggregated Ethernet Interfaces*
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39

lacp (802.3ad)

Syntax

```
lacp {
  force-up;
  port-priority
}
```

Hierarchy Level

[edit interfaces *interface-name* ether-options 802.3ad]

[edit interfaces ae*X* aggregated-ether-options]

[edit chassis aggregated-devices ethernet]

Release Information

Statement introduced in Junos OS Release 10.0 for EX Series switches.
Support for LACP link protection introduced in Junos OS Release 11.4 for EX Series switches.

Description

Configure the Link Aggregation Control Protocol (LACP) parameters for aggregated Ethernet interfaces on the global level (for all the aggregated Ethernet interfaces on the switch) or for a specific aggregated Ethernet interface.

Required Privilege Level

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

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- Understanding Aggregated Ethernet Interfaces and LACP on page 8
- *Junos OS Ethernet Interfaces Configuration Guide*

link-mode

Syntax link-mode *mode*;

Hierarchy Level [edit interfaces *interface-name*],
[edit interfaces *interface-name* ether-options],
[edit interfaces ge-*pim*/0/0 switch-options switch-port *port-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description Set the device’s link connection characteristic.

Options *mode*—Link characteristics:

- **automatic**—Link mode is negotiated. This is the default for EX Series switches.
- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.

Default: Fast Ethernet interfaces, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode. The router’s management Ethernet interface, **fxp0** or **em0**, the built-in Fast Ethernet interfaces on the FIC (M7i router), and the Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.

NOTE: On J Series ePIM Fast Ethernet interfaces, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."

NOTE:

- On EX4300 switches, the interfaces operate in full-duplex mode only.
- On EX Series switches, if **no-auto-negotiation** is specified in **[edit interfaces *interface-name* ether-options]**, you can select only **full-duplex** or **half-duplex**. If **auto-negotiation** is specified, you can select any mode.

NOTE: Member links of an aggregated Ethernet bundle must not be explicitly configured with a link mode. You must remove any such link-mode configuration before committing the aggregated Ethernet configuration.

Required Privilege Level

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

Related Documentation

- *Configuring the Link Characteristics on Ethernet Interfaces*
- *Understanding Management Ethernet Interfaces*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

link-protection

Syntax

```
link–protection {
  disable;
  (revertive |non-revertive);
}
```

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]
[edit interfaces ae*x* aggregated-ether-options *lacp*]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for **disable**, **revertive**, and **non-revertive** statements added in Junos OS Release 9.3.

Description

On the router, for aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links on the router, include the **primary** and **backup** statements at the **[edit interfaces ge-*fpc/pic/port* gigether-options 802.3ad ae*x*]** hierarchy level or the **[edit interfaces fe-*fpc/pic/port* fastether-options 802.3ad ae*x*]** hierarchy level.

On the switch, you can configure either Junos OS link protection for aggregated Ethernet interfaces or the LACP standards link protection for aggregated Ethernet interfaces.

For Junos OS link protection, specify **link-protection** at the following hierarchy levels:

- **[edit interfaces ge-*fpc/pic/port* ether-options 802.3ad ae*x*]**
- **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]**

For LACP standards link protection, specify **link-protection** at the following hierarchy levels:

- For global LACP link protection, specify at **[edit chassis aggregated-devices ethernet lacp]**
- For a specific aggregated Ethernet interface, specify at **[edit interfaces aeX aggregated-ether-options lacp]**

To disable link protection, use the **delete interface ae aggregate-ether-options link-protection** statement at the **[edit interfaces ae*x* aggregated-ether-options]** hierarchy level or the **[edit interfaces ae*x* aggregated-ether-options *lacp*]** hierarchy level.

Options

The 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 Aggregated Ethernet Link Protection on page 130

- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

link-protection-sub-group (802.3ad)

Syntax link-protection-sub-group *group-name*;

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

Release Information Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description Add an interface in an aggregated ethernet bundle into a link protection subgroup.

A link protection subgroup is created and named using the link-protection-sub-group statement in the [edit interfaces ae*x* **aggregated-ether-options**] hierarchy.

Options ***group-name***—Name of the link protection subgroup that will include this interface after this statement is entered. The link protection subgroup is named when it is created using the link-protection-sub-group statement in the [edit interfaces ae*x* **aggregated-ether-options**] hierarchy.

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

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 130
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

link-protection-sub-group (aggregated-ether-options)

Syntax

```
link-protection-sub-group group-name {
  [primary | backup];
}
```

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options]

Release Information Statement introduced in Junos OS Release 14.1X53-D10 for EX Series switches.

Description Create and name a link protection subgroup.

Link protection subgroups allow you to provide link protection to a collection of Ethernet links within a LAG bundle. If you need to provide link protection to a single link in an aggregated ethernet bundle, you do not need to configure link protection subgroups.

A link protection subgroup includes multiple links within the aggregated ethernet bundle. If one link in the primary link protection subgroup fails, traffic is forwarded over the links in the backup link protection subgroup.

Links within the aggregated ethernet bundle are added to the link protection subgroup using the link-protection-sub-group statement in the [**edit interfaces *interface-name* ether-options 802.3ad**] hierarchy.

Options ***group-name***—Creates and names the link protection subgroup. The name is created by the user.

primary—Specifies that the subgroup is the primary subgroup.

backup—Specifies that the subgroup is the backup subgroup.

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

Related Documentation

- Configuring Aggregated Ethernet Link Protection on page 130
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

link-speed (Aggregated Ethernet)

Syntax link-speed *speed*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-ether-options],
[edit interfaces interface-range *name* aggregated-sonet-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description For aggregated Ethernet interfaces only, set the required link speed.

Options *speed*—For aggregated Ethernet links, you can specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speeds:

- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **oc192**—Links are OC192 or STM64c.

Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:

- **10m**—Links are 10 Mbps.
- **100m**—Links are 100 Mbps.
- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.

Aggregated Ethernet links on T Series routers can be configured to operate at one of the following speeds:

- **100g**—Links are 100 Gbps.
- **100m**—Links are 100 Mbps.
- **10g**—Links are 10 Gbps.
- **1g**—Links are 1 Gbps.
- **40g**—Links are 40 Gbps.
- **50g**—Links are 50 Gbps.
- **80g**—Links are 80 Gbps.
- **8g**—Links are 8 Gbps.

- **mixed**—Links are of various speeds.
- **oc192**—Links are OC192.

Required Privilege Level

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

Related Documentation

- *Aggregated Ethernet Interfaces Overview*
- Configuring Aggregated Ethernet Link Speed on page 132
- *Configuring Mixed Aggregated Ethernet Links*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33

local-bias (edit interfaces ae)

Syntax

local-bias
<disable>;

Hierarchy Level

[edit interfaces ae*x* aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices.
disable option introduced in Junos OS Release 14.1X53-D25.

Description

Enable or disable local link bias for individual aggregated Ethernet interfaces.

Local link bias can be enabled or disabled globally for the entire Virtual Chassis or VCF using the **set forwarding-options local-bias** statement or **set forwarding-options local-bias disable** statement, or per LAG bundle using the **set interfaces ae*x* aggregated-ether-options** local-bias (edit interfaces ae) statement or **set interfaces ae*x* aggregated-ether-options** local-bias (edit interfaces ae) **disable** statement.

When local link bias is set at both the global and per LAG bundle levels, the per LAG bundle configuration takes precedence. For instance, if local link bias is enabled globally on the Virtual Chassis or VCF using the **set forwarding-options local-bias** statement but disabled on a LAG bundle named **ae1** using the **set interfaces ae1 aggregated-ether-options** local-bias (edit interfaces ae) **disable** statement, local link bias is disabled on the LAG bundle named **ae1**.

Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a link aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.

Default

Local link bias is disabled by default.

Options

none—Enable local link bias for the aggregated Ethernet interface.

disable—(Optional) Disable local link bias for the aggregated Ethernet interface.

Required Privilege Level

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

Related Documentation

- Configuring Local Link Bias (CLI Procedure) on page 147
- Understanding Local Link Bias on page 28

local-bias (forwarding-options)

Syntax

local-bias
<disable>;

Hierarchy Level

[edit forwarding-options]

Release Information

Statement introduced in Junos OS Release 14.1X53-D25 for EX Series switches and QFX Series devices.

Description

Enable or disable local link bias globally for all aggregated Ethernet interfaces in a Virtual Chassis or Virtual Chassis Fabric (VCF).

Local link bias can be enabled or disabled globally for the entire Virtual Chassis or VCF using the **set forwarding-options local-bias** statement or **set forwarding-options local-bias disable** statement, or per LAG bundle using the **set interfaces ae*x* aggregated-ether-options** local-bias (edit interfaces ae) statement or **set interfaces ae*x* aggregated-ether-options** local-bias (edit interfaces ae) **disable** statement.

When local link bias is set at both the global and per LAG bundle levels, the per LAG bundle configuration takes precedence. For instance, if local link bias is enabled globally on the Virtual Chassis or VCF using the **set forwarding-options local-bias** statement but disabled on a LAG bundle named **ae1** using the **set interfaces ae1 aggregated-ether-options** local-bias (edit interfaces ae) **disable** statement, local link bias is disabled on the LAG bundle named **ae1**.

Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a link aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic.

You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.

Default

Local link bias is disabled by default.

Options

none—Enable local link bias globally for the aggregated Ethernet interfaces in the Virtual Chassis or VCF.

disable—(Optional) Disable local link bias globally for the aggregated Ethernet interfaces in the Virtual Chassis or VCF.

Required Privilege Level

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

Related Documentation

- Configuring Local Link Bias (CLI Procedure) on page 147

- Understanding Local Link Bias on page 28

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

Syntax (loopback | no-loopback);

Hierarchy Level [edit interfaces *interface-name* aggregated-ether-options],
[edit interfaces *interface-name* ether-options],
[edit interfaces *interface-name* fastether-options],
[edit interfaces *interface-name* gigether-options],
[edit interfaces interface-range *name* ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.

NOTE:

- By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.
- IPv6 Neighbor Discovery Protocol (NDP) addresses are not supported on Gigabit Ethernet interfaces when loopback mode is enabled on the interface. That is, if the **loopback** statement is configured at the **[edit interfaces ge-*fpc/pic/port* gigether-options]** hierarchy level, an NDP address cannot be configured at the **[edit interfaces ge-*fpc/pic/port* unit *logical-unit-number* family inet6 address]** hierarchy level.

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

Related Documentation

- Configuring Ethernet Loopback Capability on page 115

mdi-mode

Syntax mdi-mode *mode*;

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

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

Description You must configure media dependent interface (MDI) properties for a 10-Gigabit Ethernet interface on a copper network port of an EX4550 switch to ensure that both sides of the link are compatible.

MDI refers to the IEEE standard for the interface to an unshielded twisted pair (UTP) cable. Twisted-pair Ethernet standards are such that the majority of cables can be wired "straight-through" (pin 1 to pin 1, pin 2 to pin 2 and so on), but others may need to be wired in the "crossover" form (receive to transmit and transmit to receive).

For most ports, the switch can automatically detect the required connection type and can therefore configure the interface appropriately. However, the switch cannot automatically detect whether the connection type of a 10-Gigabit Ethenet interface on a copper network port is straight-through or crossover.

Therefore, you must set the MDI properties of the local interface of a 10-Gigabit Ethernet interface on a copper network port to ensure that it will work correctly with the other side of the link. When you set this configuration on an interface, you must also disable auto-negotiation and set the speed to **100m**.

NOTE: **This configuration does not apply to management Ethernet or console interfaces and it does not apply to 1-Gigabit copper ports.**

The proper setting depends both on the type of cable and the setting that is being used on the other side of the link:

- For crossover cables—Set the polarity to match the other side of the link. Specify **mdi** for the switch interface if **mdi** is being used on the other side of the link; specify **mdix** if **mdix** is being used on the other side of the link.
- For straight cables—Set the polarity to be the opposite of the other side link. Specify **mdi** for the switch interface if **mdix** is being used on the other side of the link; specify **mdix** if **mdi** is being used on the other side of the link.

Options One of the following modes:

auto—Set the MDI properties to automatic. This setting should *not* be used with 10-Gigabit Ethernet interfaces on a copper network port of an EX4550 switch.

mdi—Set the MDI properties of the interface to straight through mode. The selection of the mode depends on whether crossover or straight cables are being used and on the setting used on the other side of the link.

mdix—Set the MDI properties of the interface to crossover mode. The selection of the mode depends on whether crossover or straight cables are being used and on the setting used on the other side of the link.

Required Privilege Level

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

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- EX Series Switches Interfaces Overview on page 3
- *Junos OS Ethernet Interfaces Configuration Guide*

media-type (Dual-Purpose Uplink Ports)

Syntax

media-type (copper | fiber);

Hierarchy Level

[edit interfaces *interface-name*]

Release Information

Statement introduced in Junos OS Release 11.3 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX1000 Universal Access Routers.

Description

(EX2200-C switch only) Configure the media type for a dual-purpose uplink port (one RJ-45 port and one SFP port) on an EX2200 switch. If you use the media-type for a dual-purpose uplink port, the alternate media type cannot be used with the port.

(ACX1000 routers only) Select the media type (copper or fiber) for the 1-Gigabit Ethernet interfaces.

Default

When **media-type** is not set, the port accepts either type of connection. The media type is fiber if a transceiver is installed in the SFP connection. If no transceiver is installed, the media type is copper.

Options

copper—The dual-purpose uplink port accepts only a 10/100/1000BASE-T copper connection.

fiber—The dual-purpose uplink port accepts only an SFP fiber connection.

Required Privilege Level

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

Related Documentation

- Configuring the Media Type on Dual-Purpose Uplink Ports (CLI Procedure) on page 141

member (Interface Ranges)

Syntax member *interface-name*;

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

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.

Description Specify the name of the member interface belonging to an interface range on the EX Series switch.

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

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

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Understanding Interface Ranges on EX Series Switches on page 17
- *Understanding Interface Ranges on EX Series Switches*
- EX Series Switches Interfaces Overview on page 3
- *Junos OS Interfaces Fundamentals Configuration Guide*

member-range

Syntax member-range *starting-interface-name* to *ending-interface-name*;

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

Release Information Statement introduced in Junos OS Release 10.0 for EX Series switches.

Description Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.

Options **Range:** *Starting interface-name* to *ending interface-name*—The name of the first member and the name of the last member in the interface sequence.

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

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Understanding Interface Ranges on EX Series Switches on page 17
- *Understanding Interface Ranges on EX Series Switches*
- EX Series Switches Interfaces Overview on page 3
- *Junos OS Interfaces Fundamentals Configuration Guide*

members

Syntax members [(all | *names* | *vlan-ids*)];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family ethernet-switching vlan]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement updated with enhanced **?** (CLI completion feature) functionality in Junos OS Release 9.5 for EX Series switches.

Description For trunk interfaces, configure the VLANs that can carry traffic.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after **vlan** or **vlans** in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command **set vlans id vlan-id ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum.

On an EX Series switch that runs Junos OS that does not support the Enhanced Layer 2 Software (ELS) configuration style, the maximum number of VLAN members allowed on the switch is 8 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 8). If the switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (**eswd**) due to memory allocation failure.

On an EX Series switch that runs Junos OS that supports ELS, the maximum number of VLAN members allowed on the switch is 24 times the maximum number of VLANs the switch supports (vmember limit = vlan max * 24). If the configuration of one of these switches exceeds the recommended VLAN member maximum, a warning message appears in the system log (syslog).

Options **all**—Specifies that this trunk interface is a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.

NOTE: **Since VLAN members are limited, specifying all could cause the number of VLAN members to exceed the limit at some point.**

names—Name of one or more VLANs. VLAN IDs are applied automatically in this case.

NOTE: **all cannot be a VLAN name.**

vlan-ids—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, **10–20** or **10–20 23 27–30**.

NOTE: **Each configured VLAN must have a specified VLAN ID to successfully commit the configuration; otherwise, the configuration commit fails.**

Required Privilege Level

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

Related Documentation

- *show ethernet-switching interfaces*
- *show ethernet-switching interface*
- *show vlans*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch*
- *Example: Connecting an Access Switch to a Distribution Switch*
- *Example: Connecting Access Switches to a Distribution Switch*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Creating a Series of Tagged VLANs (CLI Procedure)*
- *Understanding Bridging and VLANs on EX Series Switches*
- *Junos OS Ethernet Interfaces Configuration Guide*

minimum-links

Syntax minimum-links *number*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options],
[edit interfaces ae*x* aggregated-sonet-options],
[edit interfaces *interface-name* mlfr-uni-nni-bundle-options],
[edit interfaces *interface-name* unit *logical-unit-number*],
[edit interfaces interface-range *range* aggregated-ether-options],
[edit interfaces interface-range *range* aggregated-sonet-options],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.

Options *number*—Number of links.

Range: On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for minimum-links number is 1 through 64. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.On all other routers and on EX Series switches, other than EX8200 switches, the range of valid values for minimum-links number is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled upOn EX8200 switches, the range of valid values for minimum-links number is 1 through 12. When the maximum value (12) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

Default: 1

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

Related Documentation

- Configuring Aggregated Ethernet Minimum Links on page 133
- *Configuring Aggregated SONET/SDH Interfaces*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- *Junos OS Services Interfaces Library for Routing Devices*

mtu

Syntax

mtu *bytes*;

Hierarchy Level

[edit interfaces *interface-name*],
[edit interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit interfaces interface-range *name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*],
[edit logical-systems *logical-system-name* protocols l2circuit local-switching interface *interface-name* backup-neighbor *address*],
[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols l2circuit neighbor *address* interface *interface-name* backup-neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols vpls],
[edit protocols l2circuit local-switching interface *interface-name* backup-neighbor *address*],
[edit protocols l2circuit neighbor *address* interface *interface-name*]
[edit protocols l2circuit neighbor *address* interface *interface-name* backup-neighbor *address*],
[edit routing-instances *routing-instance-name* protocols l2vpn interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols l2vpn site *site-name*],
[edit routing-instances *routing-instance-name* protocols vpls]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Support at the **[set interfaces interface-name unit logical-unit-number family ccc]** hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.

Description

Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

To route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named irb or vlan, respectively).

CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.

NOTE:

The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the **flexible-vlan-tagging** statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
- In case of Layer 2 IFL configured with the **vlan-tagging** statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.

NOTE:

- If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.
- Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.
- On ACX Series routers, you can configure the protocol MTU by including the **mtu** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** or **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level.
  - If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
  - If you are configuring the protocol MTU for both **inet** and **inet6** families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for **inet** and **inet6** families that are configured on the same logical interface.

For more information about configuring MTU for specific interfaces and router or switch combinations, see “Configuring the Media MTU” on page 90.

Options *bytes*—MTU size.

Range: 256 through 9192 bytes, 256 through 9216 (EX Series switch interfaces), 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers)

Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)

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

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring the Media MTU on page 90
- *Configuring the MTU for Layer 2 Interfaces*
- Setting the Protocol MTU on page 103

multicast-loadbalance

Syntax

```
multicast-loadbalance {
  disable;
  hash-mode mode;
}
```

Hierarchy Level [edit chassis]

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

Description (EX8200 switches only) Enable multicast load balancing to evenly distribute the Layer 3 routed multicast traffic on aggregated 10-gigabit Ethernet interfaces. When multicast load balancing is enabled on the EX8200 switch, it takes effect on aggregated 10-gigabit Ethernet interfaces that are currently configured as well as those that are configured later.

Crc-sgip mode is the default. If the default mode is implemented and the Layer 3 routed multicast traffic on the aggregated interfaces is not evenly distributed, you can try the other hashing algorithms. Use the algorithm that best balances your Layer 3 routed multicast traffic.

NOTE: **If you change the hashing algorithm when multicast load balancing is disabled, the new algorithm takes effect after you reenable multicast load balancing.**

Default Multicast load balancing is enabled.

Options **hash-mode *mode***—Specify one of these hashing algorithms:

balanced—Use a round-robin method to determine the aggregated Ethernet links with the least amount of traffic.

crc-gip —Use multicast packets' group IP address as a key to a CRC algorithm to get the hash index. The hashing algorithm used is CRC.

crc-sgip —Use multicast packets' source IP address and group IP address for hashing. The hashing algorithm used is CRC.

crc-sip —Use multicast packets' source IP address for hashing. The hashing algorithm used is CRC.

simple-gip—Use multicast packets' group IP address for hashing. The hashing algorithm used is XOR.

simple-sgip—Use multicast packets' GIP bits to calculate the hash index. The hashing algorithm used is XOR.

simple-sip—Use multicast packets' source IP address bits for hashing. The hashing algorithm used is XOR.

The remaining statement is explained separately.

Required Privilege Level

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

Related Documentation

- Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches on page 59
- Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Links on EX8200 Switches (CLI Procedure) on page 144

native-vlan-id

Syntax

native-vlan-id *vlan-id*;

Hierarchy Level

[edit interfaces *interface-name* unit 0 family ethernet-switching]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure the VLAN identifier to associate with untagged packets received on the interface.

Options

vlan-id—Numeric identifier of the VLAN.

Range: 0 through 4095

Required Privilege Level

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

Related Documentation

- *show vlans*
- *show ethernet-switching interfaces*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- *Understanding Bridging and VLANs on EX Series Switches*
- *Junos OS Ethernet Interfaces Configuration Guide*

nd6-stale-time

Syntax nd6-stale-time *seconds*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6]

Release Information Statement introduced in Junos OS Release 11.1.
Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description Set the stale timer for IPv6 neighbor reachability confirmation. Reachability of the IPv6 neighbors is confirmed only after the stale timer has expired. For example, by setting the stale timer to 180 seconds, users can specify that IPv6 neighbor reachability be confirmed every 180 seconds.

NOTE: When the Routing Engine sends a control packet to an IPv6 neighbor, the stale timer is the maximum interval in which neighbor reachability is confirmed. In such cases, IPv6 neighbor reachability is confirmed before the stale timer expires.

Default Default is 20 minutes (1200 seconds)

Options *seconds*—Duration in seconds.

Range: 1 to 1200

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

Related Documentation
- *IPv6 Neighbor Discovery Overview*
- *show ipv6 neighbors*

no-redirects

Syntax no-redirects;

Hierarchy Level [edit 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 9.0 for EX Series switches.

Description Do not send protocol redirect messages on the interface.

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the **[edit system]** hierarchy level.

Default Interfaces send protocol redirect messages.

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

Related Documentation

- Disabling the Transmission of Redirect Messages on an Interface on page 118
- *Junos OS Administration Library for Routing Devices*

non-revertive (Chassis)

Syntax non-revertive;

Hierarchy Level [edit chassis aggregated-devices ethernet lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Disable the ability to switch to a better priority link (if one is available) once a link is established as active and a collection or distribution is enabled.

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

Related Documentation

- *Configuring Junos OS for Supporting Aggregated Devices*
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

non-revertive (Interfaces)

Syntax non-revertive;

Hierarchy Level [edit interfaces ae*X* aggregated-ether-options lacp link-protection]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.

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

Related Documentation

- **link-protection on page 232**
- Configuring Aggregated Ethernet Link Protection on page 130
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

optics-options

Syntax

```
optics-options {
  alarm low-light-alarm {
    (link-down | syslog);
  }
  tx-power dbm;
  warning low-light-warning {
    (link-down | syslog);
  }
  wavelength nm;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.
alarm option and **warning** options introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement and **tx-power** option introduced in Junos OS Release 13.2 for PTX Series routers.

Description For 10-Gigabit Ethernet or 100-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces only, configure full C-band International Telecommunication Union (ITU)-Grid tunable optics.

Options 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

- *Ethernet DWDM Interface Wavelength Overview*
- *100-Gigabit Ethernet OTN Options Configuration Overview*

periodic

Syntax periodic *interval*;

Hierarchy Level [edit interfaces ae*x* aggregated-ether-options lacp],
[edit interfaces interface-range *name* aggregated-ether-options lacp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.

Options ***interval***—Interval for periodic transmission of LACP packets.

- **fast**—Transmit packets every second.
- **slow**—Transmit packets every 30 seconds.

Default: fast

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

Related Documentation

- *Configuring LACP for Aggregated Ethernet Interfaces*
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33

pic

Syntax

```
pic pic-number {
  sfpplus {
    pic-mode mode;
    }
  tunnel-port port-number tunnel-services;
}
```

Hierarchy Level [edit chassis fpc *slot*]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Enable the specified port of the SFP+ uplink module to perform in the operating mode specified by pic-mode. The port is indicated by a Physical Interface Card (PIC) number.

For generic routing encapsulation (GRE) tunneling, use the **pic** statement with the tunnel-port statement to specify the number of the port on the switch that you want to convert to a GRE tunnel port.

Options *pic-number*—Number of the PIC. For built-in ports on EX3200, EX4200, EX4500, and EX4550 switches and on EX6200 and EX8200 switches, the PIC number is 0. For uplink ports on EX3200 and EX4200 switches, the PIC number is 1. On EX4500 switches, the PIC number is 1 for ports on the left-hand uplink module and 2 for ports on the right-hand uplink module. On EX4550 switches, the PIC number is 1 for ports in the expansion module installed in the module slot on the front panel of the switch and 2 for those in the expansion module installed in the module slot on the rear panel of the switch.

NOTE: Configuration of an uplink port on an EX4500 or EX4550 switch as a GRE tunnel port is not supported

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

- Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure) on page 140
- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) on page 142

pic-mode

Syntax pic-mode *mode*;

Hierarchy Level [edit chassis fpc *slot* pic *pic-number* sfpplus]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Configure the operating mode for the specified port on the SFP+ uplink module on an EX3200 or EX4200 switch.

On a switch using the SFP+ or SFP+ MACSec uplink module, the **pic-mode** setting defines the speeds of the SFP+ interfaces. When the PIC mode is set to **10g** on the SFP+ or SFP+ MACSec uplink module, uplink ports 0 and 2 support MACSec at 10-Gbps speeds while ports 1 and 3 cannot be used to send any traffic. When the PIC mode is set to **1g**, all four SFP+ ports on the uplink module support MACSec at 1-Gbps speeds.

Options *mode*—Operating mode of the SFP+ uplink module:

- **1G**—1-gigabit operating mode
- **10G**—10-gigabit operating mode

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

Related Documentation

- Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure) on page 140

port-priority

Syntax port-priority *priority*;

Hierarchy Level [edit interfaces *interface-name* gigether-options 802.3ad lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Define LACP port priority at the interface level.

Options *priority*—Priority for being elected to be the active port and both collect and distribute traffic. A smaller value indicates a higher priority for being elected.

Range: 1 through 255

Default: 127

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

Related Documentation

- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- *Configuring Aggregated Ethernet LACP*

port-mode

Syntax port-mode *mode*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family ethernet-switching]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure whether an interface on the switch operates in access, tagged-access, or trunk mode.

Default All switch interfaces are in access mode.

Options *mode*—Operating mode for an interface can be one of the following:

- **access**—In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to single network devices such as PCs, printers, IP telephones, and IP cameras.
- **tagged-access**—In this mode, the interface can accept tagged packets from one access device. Tagged-access interfaces typically connect to servers running Virtual machines using VEPA technology.
- **trunk**—In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN.

NOTE: The number of VLANs supported per switch varies for each model. Use the configuration-mode command **set vlans id vlan-id ?** to determine the maximum number of VLANs allowed on a switch. You cannot exceed this VLAN limit because each VLAN is assigned an ID number when it is created. You can, however, exceed the recommended VLAN member maximum. To determine the maximum number of VLAN members allowed on a switch, multiply the VLAN maximum for the switch times 8 (vmember limit = vlan max * 8).

If a switch configuration exceeds the recommended VLAN member maximum, you see a warning message when you commit the configuration. If you ignore the warning and commit such a configuration, the configuration succeeds but you run the risk of crashing the Ethernet switching process (eswd) due to memory allocation failure.

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

Related Documentation

- *Example: Connecting an Access Switch to a Distribution Switch*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66

- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

preferred

Syntax preferred;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.

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

Default The lowest-numbered address on the subnet is the preferred address.

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

Related Documentation
- Configuring the Interface Address on page 84

primary (Address on Interface)

Syntax primary;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.

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

Default For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.

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

Related Documentation

- Configuring the Interface Address on page 84

proxy-arp

Syntax
proxy-arp (restricted | unrestricted);

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

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
restricted added in Junos OS Release 10.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for the QFX Series.

Description
For Ethernet interfaces only, configure the router or switch to respond to any ARP request, as long as the router or switch has an active route to the ARP request's target address.

NOTE: You must configure the IP address and the inet family for the interface when you enable proxy ARP.

Default
Proxy ARP is not enabled. The router or switch responds to an ARP request only if the destination IP address is its own.

Options

- **none**—The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.
- **restricted**—(Optional) The router or switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are in the same subnet. The router or switch must also have a route to the target IP address.
- **unrestricted**—(Optional) The router or switch responds to any ARP request for a local or remote address if the router or switch has a route to the target IP address.

Default: unrestricted

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

Related Documentation

- Configuring Restricted and Unrestricted Proxy ARP on page 119
- *Configuring Proxy ARP (CLI Procedure)*
- *Configuring Proxy ARP (CLI Procedure)*
- *Example: Configuring Proxy ARP on an EX Series Switch*
- Configuring Gratuitous ARP on page 116

rpf-check

Syntax rpf-check;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit interfaces *interface-name* unit *logical-unit-number* family inet6]

Release Information Statement introduced in Junos OS Release 9.3 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.

Description On EX3200 and EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on all ingress interfaces.

On EX4300 switches, enable a reverse-path forwarding (RPF) check on unicast traffic, including ECMP packets, on all ingress interfaces.

On EX8200 and EX6200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interfaces.

On QFX Series switches, enable an RPF check on unicast traffic (except ECMP packets) on the selected ingress interfaces.

Default Unicast RPF is disabled on all interfaces.

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

Related Documentation

- Example: Configuring Unicast RPF on an EX Series Switch on page 51
- Configuring Unicast RPF (CLI Procedure) on page 135
- Disabling Unicast RPF (CLI Procedure) on page 137
- Understanding Unicast RPF on page 12

sfpplus

Syntax

```
sfpplus {
  pic-mode mode;
}
```

Hierarchy Level [edit chassis fpc *slot* pic *pic-number*]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Configure the operating mode for the specified port on the SFP+ uplink module on the EX3200 or EX4200 switch.

The remaining statement is explained separately.

Default By default, the SFP+ uplink module operates in the 10-gigabit mode and supports SFP+ transceivers.

NOTE: The SFP+ uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.

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

Related Documentation

- Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure) on page 140

source

Syntax source *source-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* tunnel]

Release Information 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.

Description Specify the source address of the tunnel.

Default If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.

Options ***source-address***—Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.

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

Related Documentation

- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) on page 142

speed

Syntax speed (auto-negotiation | *speed*) ;

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

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the interface's speed.

Default If the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled, the **auto-negotiation** option is enabled by default.

Options

- **auto-negotiation**—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled.
- ***speed***—Specify the interface speed. If the auto-negotiation statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is disabled, you must specify a specific value. This value sets the speed that is used on the link. If the **auto-negotiation** statement is enabled, you might want to configure a specific speed value to advertise the desired speed to the remote end.
  - **10m**—10 Mbps
  - **100m**—100 Mbps
  - **1g**—1 Gbps

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

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- *Junos OS Ethernet Interfaces Configuration Guide*

system-priority

Syntax system-priority *priority*;

Hierarchy Level [edit chassis aggregated-devices ethernet lacp]

Release Information Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Define LACP system priority for aggregated Ethernet interfaces at the global (chassis) level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

Options *priority*—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

Range: 0 through 65535

Default: 127

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

Related Documentation

- *Configuring Junos OS for Supporting Aggregated Devices*
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

system-priority

Syntax

system-priority *priority*;

Hierarchy Level

[edit interfaces ae*X* aggregated-ether-options lacp]

Release Information

Statement introduced in Junos OS Release 9.3.
Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description

Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global **[edit chassis]** hierarchy level.

The device with the lower system priority value determines which links between LACP partner devices are active and which are in standby for each LACP group. The device on the controlling end of the link uses port priorities to determine which ports are bundled into the aggregated bundle and which ports are put in standby mode. Port priorities on the other device (the noncontrolling end of the link) are ignored. In priority comparisons, numerically lower values have higher priority. Therefore, the system with the numerically lower value (higher priority value) for LACP system priority becomes the controlling system. If both devices have the same LACP system priority (for example, they are both configured with the default setting of 127), the device MAC address determines which switch is in control.

Options

priority—Priority for the aggregated Ethernet system. A smaller value indicates a higher priority.

Range: 0 through 65535

Default: 127

Required Privilege Level

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

Related Documentation

- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126

targeted-broadcast

Syntax targeted-broadcast;

Hierarchy Level [edit interfaces ge-*chassis/slot/port* unit *logical-unit-number* family inet]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Enable IP directed broadcast on a specified subnet.

Default IP directed broadcast is disabled.

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

Related Documentation

- Example: Configuring IP Directed Broadcast on an EX Series Switch on page 56
- Configuring IP Directed Broadcast (CLI Procedure) on page 137
- Understanding IP Directed Broadcast for EX Series Switches on page 16

traceoptions (Individual Interfaces)

Syntax

```
traceoptions {
  file filename <files name> <size size> <world-readable | no-world-readable>;
  flag flag;
  match;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description Define tracing operations for individual interfaces.

To specify more than one tracing operation, include multiple **flag** statements.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** file in the directory **/var/log/dcd**.

Default If you do not include this statement, no interface-specific tracing operations are performed.

Options

file name—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log/dcd**. By default, interface process tracing output is placed in the file **files number**—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

match—(Optional) Regular expression for lines to be traced.

no-world-readable—(Optional) Prevent any user from reading the log file.

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

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

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options.

- **all**—All interface tracing operations
- **event**—Interface events
- **ipc**—Interface interprocess communication (IPC) messages

- **media**—Interface media changes
- **q921**—Trace ISDN Q.921 frames
- **q931**—Trace ISDN Q.931 frames

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

Related Documentation
- *Tracing Operations of an Individual Router Interface*

traceoptions (Interface Process)

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
  no-remote-trace;
}
```

Hierarchy Level [edit interfaces]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Define tracing operations for the interface process (dcd).

Default If you do not include this statement, no interface-specific tracing operations are performed.

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

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, interface process tracing output is placed in the file **dcd**.

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

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**
- **change-events**—Log changes that produce configuration events
- **config-states**—Log the configuration state machine changes
- **kernel**—Log configuration IPC messages to kernel
- **kernel-detail**—Log details of configuration messages to kernel

no-world-readable—(Optional) Disallow any user to read the log file.

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

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify kilobytes, ***x*m** to specify megabytes, or ***x*g** to specify gigabytes

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

Default: 1 MB

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

match ***regex***—(Optional) Refine the output to include only those lines that match the given regular expression.

Required Privilege Level

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

Related Documentation

- Tracing Operations of the Interface Process on page 139

traps

Syntax

(traps | no-traps);

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description

Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

Required Privilege Level

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

Related Documentation

- Enabling or Disabling SNMP Notifications on Physical Interfaces on page 120
- Enabling or Disabling SNMP Notifications on Logical Interfaces on page 120

ttl

Syntax ttl *value*;

Hierarchy Level [edit interfaces *interface-name* unit *number* tunnel]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.1 for EX Series switches.

Description Set the time-to-live value bit in the header of the outer IP packet.

Options ***value***—Time-to-live value.

Range: 0 through 255

Default: 64

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

Related Documentation

- *Tunnel Properties*
- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) on page 142

tunnel

Syntax

```
tunnel {
  destination destination-address;
  source source-address;
  ttl number; (not supported on QFX and OCX Series switches)
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information 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.

Description 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.

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 Generic Routing Encapsulation Tunneling (CLI Procedure) on page 142

tunnel-port

Syntax tunnel-port *port-number* tunnel-services;

Hierarchy Level [edit chassis fpc *slot* pic *pic-number*]

Release Information 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.

Description Configure the port number for generic routing encapsulation (GRE) tunneling.

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

Related Documentation

- Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) on page 142

unit

Syntax

```
unit logical-unit-number {
  accounting-profile name;
  bandwidth rate;
  description text;
  disable;
  family family-name {...}
  proxy-arp (restricted | unrestricted);
  (traps | no-traps);
  vlan-id (VLAN Tagging and Layer 3 Subinterfaces) vlan-id-number;
}
```

Hierarchy Level

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

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure a logical interface on the physical device. 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.

Required Privilege Level

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

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- EX Series Switches Interfaces Overview on page 3
- *Junos OS Ethernet Interfaces Configuration Guide*

vlan (802.1Q Tagging)

Syntax

```
vlan {
  members [(all | names | vlan-ids)];
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family ethernet-switching]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Bind an 802.1Q VLAN tag ID to a logical interface.

The remaining statement is explained separately.

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

Related Documentation

- *show ethernet-switching interfaces*
- *show ethernet-switching interface*
- *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*
- *Configuring Routed VLAN Interfaces (CLI Procedure)*
- *Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)*
- *Understanding Bridging and VLANs on EX Series Switches*
- *Junos OS Ethernet Interfaces Configuration Guide*

vlan-id (VLAN Tagging and Layer 3 Subinterfaces)

Syntax vlan-id *vlan-id-number*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2 for EX Series switches.

Description Bind an 802.1Q VLAN tag ID to a logical interface.

NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

Options *vlan-id-number*—A valid VLAN identifier.

Range: 1 through 4094

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

Related Documentation

- vlan-tagging on page 280
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 44
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 135
- *Configuring Q-in-Q Tunneling (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

vlan-tagging

Syntax vlan-tagging;

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2 for PTX Series Routers.
Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series.

Description For Fast Ethernet and Gigabit Ethernet interfaces, aggregated Ethernet interfaces configured for VPLS, and pseudowire subscriber interfaces, enable the reception and transmission of 802.1Q VLAN-tagged frames on the interface.

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

Related Documentation

- *Example: Configuring BGP Autodiscovery for LDP VPLS*
- Configuring Tagged Aggregated Ethernet Interfaces on page 134
- *Configuring Interfaces for VPLS Routing*
- *Enabling VLAN Tagging*
- 802.1Q VLANs Overview on page 19

warning

Syntax

```
warning low-light-warning {
   (link-down | syslog);
}
```

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 12.1 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series routers.

Description Specifies the action to take if the receiving optics signal is below the optics low-light warning threshold.

Options **link-down**—Drop the 10-Gigabit Ethernet link and marks link as down.

syslog—Write the optics information to the system log.

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

Related Documentation

- *Configuring 10-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning*
- optics-options on page 255
- *100-Gigabit Ethernet OTN Options Configuration Overview*

wavelength

Syntax wavelength *nm*;

Hierarchy Level [edit interfaces *interface-name* optics-options]

Release Information 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 PTX Series routers.

Description For 10-Gigabit or 100-Gigabit Ethernet DWDM interfaces only, configure full C-band ITU-Grid tunable optics.

Options *nm*—Wavelength value. It can be one of the following:

NOTE: All values are displayed. However, if you configure a value that is not supported by the device, an error message is displayed and the device is not tuned to the specified wavelength.

- **1528.38**—1528.38 nanometers (nm), corresponds to a 50-GHz grid
- **1528.77**—1528.77 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.16**—1529.16 nm, corresponds to a 50-GHz grid
- **1529.55**—1529.55 nm, corresponds to 50-GHz and 100-GHz grids
- **1529.94**—1529.94 nm, corresponds to a 50-GHz grid
- **1530.33**—1530.33 nm, corresponds to 50-GHz and 100-GHz grids
- **1530.72**—1530.72 nm, corresponds to a 50-GHz grid
- **1531.12**—1531.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1531.51**—1531.51 nm, corresponds to a 50-GHz grid
- **1531.90**—1531.90 nm, corresponds to 50-GHz and 100-GHz grids
- **1532.29**—1532.29 nm, corresponds to a 50-GHz grid
- **1532.68**—1532.68 nm, corresponds to 50-GHz and 100-GHz grids
- **1533.07**—1533.07 nm, corresponds to a 50-GHz grid
- **1533.47**—1533.47 nm, corresponds to 50-GHz and 100-GHz grids
- **1533.86**—1533.86 nm, corresponds to a 50-GHz grid
- **1534.25**—1534.25 nm, corresponds to 50-GHz and 100-GHz grids
- **1534.64**—1534.64 nm, corresponds to a 50-GHz grid
- **1535.04**—1535.04 nm, corresponds to 50-GHz and 100-GHz grids

- **1535.43**—1535.43 nm, corresponds to a 50-GHz grid
- **1535.82**—1535.82 nm, corresponds to 50-GHz and 100-GHz grids
- **1536.22**—1536.22 nm, corresponds to a 50-GHz grid
- **1536.61**—1536.61 nm, corresponds to 50-GHz and 100-GHz grids
- **1537.00**—1537.00 nm, corresponds to a 50-GHz grid
- **1537.40**—1537.40 nm, corresponds to 50-GHz and 100-GHz grids
- **1537.79**—1537.79 nm, corresponds to a 50-GHz grid
- **1538.19**—1538.19 nm, corresponds to 50-GHz and 100-GHz grids
- **1538.58**—1538.58 nm, corresponds to a 50-GHz grid
- **1538.98**—1538.98 nm, corresponds to 50-GHz and 100-GHz grids
- **1539.37**—1539.37 nm, corresponds to a 50-GHz grid
- **1539.77**—1539.77 nm, corresponds to 50-GHz and 100-GHz grids
- **1540.16**—1540.16 nm, corresponds to a 50-GHz grid
- **1540.56**—1540.56 nm, corresponds to 50-GHz and 100-GHz grids
- **1540.95**—1540.95 nm, corresponds to a 50-GHz grid
- **1541.35**—1541.35 nm, corresponds to 50-GHz and 100-GHz grids
- **1541.75**—1541.75 nm, corresponds to a 50-GHz grid
- **1542.14**—1542.14 nm, corresponds to 50-GHz and 100-GHz grids
- **1542.54**—1542.54 nm, corresponds to a 50-GHz grid
- **1542.94**—1542.94 nm, corresponds to 50-GHz and 100-GHz grids
- **1543.33**—1543.33 nm, corresponds to a 50-GHz grid
- **1543.73**—1543.73 nm, corresponds to 50-GHz and 100-GHz grids
- **1544.13**—1544.13 nm, corresponds to a 50-GHz grid
- **1544.53**—1544.53 nm, corresponds to 50-GHz and 100-GHz grids
- **1544.92**—1544.92 nm, corresponds to a 50-GHz grid
- **1545.32**—1545.32 nm, corresponds to 50-GHz and 100-GHz grids
- **1545.72**—1545.72 nm, corresponds to a 50-GHz grid
- **1546.12**—1546.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1546.52**—1546.52 nm, corresponds to a 50-GHz grid
- **1546.92**—1546.92 nm, corresponds to 50-GHz and 100-GHz grids
- **1547.32**—1547.32 nm, corresponds to a 50-GHz grid
- **1547.72**—1547.72 nm, corresponds to 50-GHz and 100-GHz grids
- **1548.11**—1548.11 nm, corresponds to a 50-GHz grid

- **1548.51**—1548.51 nm, corresponds to 50-GHz and 100-GHz grids
- **1548.91**—1548.91 nm, corresponds to a 50-GHz grid
- **1549.32**—1549.32 nm, corresponds to 50-GHz and 100-GHz grids
- **1549.72**—1549.72 nm, corresponds to a 50-GHz grid
- **1550.12**—1550.12 nm, corresponds to 50-GHz and 100-GHz grids
- **1550.52**—1550.52 nm, corresponds to a 50-GHz grid
- **1550.92**—1550.92 nm, corresponds to 50-GHz and 100-GHz grids
- **1551.32**—1551.32 nm, corresponds to a 50-GHz grid
- **1551.72**—1551.72 nm, corresponds to 50-GHz and 100-GHz grids
- **1552.12**—1552.12 nm, corresponds to a 50-GHz grid
- **1552.52**—1552.52 nm, corresponds to 50-GHz and 100-GHz grids
- **1552.93**—1552.93 nm, corresponds to a 50-GHz grid
- **1553.33**—1554.33 nm, corresponds to 50-GHz and 100-GHz grids
- **1553.73**—1554.73 nm, corresponds to a 50-GHz grid
- **1554.13**—1554.13 nm, corresponds to 50-GHz and 100-GHz grids
- **1554.54**—1554.54 nm, corresponds to a 50-GHz grid
- **1554.94**—1554.94 nm, corresponds to 50-GHz and 100-GHz grids
- **1555.34**—1555.34 nm, corresponds to a 50-GHz grid
- **1555.75**—1555.75 nm, corresponds to 50-GHz and 100-GHz grids
- **1556.15**—1556.15 nm, corresponds to a 50-GHz grid
- **1556.55**—1556.55 nm, corresponds to 50-GHz and 100-GHz grids
- **1556.96**—1556.96 nm, corresponds to a 50-GHz grid
- **1557.36**—1557.36 nm, corresponds to 50-GHz and 100-GHz grids
- **1557.77**—1557.77 nm, corresponds to a 50-GHz grid
- **1558.17**—1558.17 nm, corresponds to 50-GHz and 100-GHz grids
- **1558.58**—1558.58 nm, corresponds to a 50-GHz grid
- **1558.98**—1558.98 nm, corresponds to 50-GHz and 100-GHz grids
- **1559.39**—1559.39 nm, corresponds to a 50-GHz grid
- **1559.79**—1559.79 nm, corresponds to 50-GHz and 100-GHz grids
- **1560.20**—1560.20 nm, corresponds to a 50-GHz grid
- **1560.61**—1560.61 nm, corresponds to 50-GHz and 100-GHz grids
- **1561.01**—1561.01 nm, corresponds to a 50-GHz grid
- **1561.42**—1561.42 nm, corresponds to 50-GHz and 100-GHz grids

- **1561.83**—1561.83 nm, corresponds to a 50-GHz grid
- **1562.23**—1562.23 nm, corresponds to 50-GHz and 100-GHz grids
- **1562.64**—1562.64 nm, corresponds to a 50-GHz grid
- **1563.05**—1563.05 nm, corresponds to 50-GHz and 100-GHz grids
- **1563.45**—1563.45 nm, corresponds to a 50-GHz grid
- **1563.86**—1563.86 nm, corresponds to 50-GHz and 100-GHz grids
- **1564.27**—1564.27 nm, corresponds to a 50-GHz grid
- **1564.68**—1564.68 nm, corresponds to 50-GHz and 100-GHz grids
- **1565.09**—1565.09 nm, corresponds to a 50-GHz grid
- **1565.50**—1565.50 nm, corresponds to 50-GHz and 100-GHz grids
- **1565.90**—1565.90 nm, corresponds to a 50-GHz grid
- **1566.31**—1566.31 nm, corresponds to 50-GHz and 100-GHz grids
- **1566.72**—1566.72 nm, corresponds to a 50-GHz grid
- **1567.13**—1567.13 nm, corresponds to 50-GHz and 100-GHz grids
- **1567.54**—1567.54 nm, corresponds to a 50-GHz grid
- **1567.95**—1567.95 nm, corresponds to 50-GHz and 100-GHz grids
- **1568.36**—1568.36 nm, corresponds to a 50-GHz grid
- **1568.77**—1568.77 nm, corresponds to 50-GHz and 100-GHz grids

Default: **1550.12**—1550.12 nm, corresponds to 50-GHz and 100-GHz grids

Required Privilege Level

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

Related Documentation

- *Ethernet DWDM Interface Wavelength Overview*
- *Configuring the 10-Gigabit or 100-Gigabit Ethernet DWDM Interface Wavelength*
- *show interfaces diagnostics optics (Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet)*

PART 3

Administration

CHAPTER 5

Routine Monitoring

Monitoring Interface Status and Traffic

Purpose

NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to view interface status or to monitor interface bandwidth utilization and traffic statistics on the EX Series switches.

The J-Web interface monitors interface bandwidth utilization and plots real-time charts to display input and output rates in bytes per second. In addition, the Interface monitoring page displays input and output packet counters and error counters in the form of charts.

Alternatively, you can enter the **show** commands in the CLI to view interface status and traffic statistics.

NOTE: For logical interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic.

NOTE: EX Series switches do not support the collection and reporting of IPv6 transit statistics. Therefore, the **IPv6 transit statistics** field in the **show interfaces** commands displays all values as **0**.

Action To view general interface information in the J-Web interface such as available interfaces, select **Monitor** > **Interfaces**. Click any interface to view details about its status.

To set up interface monitoring for Virtual Chassis and EX8200 switches, select a member from the **Port for Member** list. Details such as the admin status and link status are displayed in the table. For an EX8200 Virtual Chassis setup, select the member, **FPC**, and the required interface.

NOTE: By default, the details of the first member in the **FPC** list is displayed. In an EX8200 Virtual Chassis setup, details of the first member and the first **FPC** is displayed.

You have the following options:

- **Start/Stop**—Starts or stops monitoring the selected interface.
- **Show Graph**—Displays input and output packet counters and error counters in the form of charts. Click the pop-up icon to view the graph in a separate window.
- **Details**—Displays interface information such as general details, traffic statistics, I/O errors, CoS counters, and Ethernet statistics.
- **Refresh Interval (sec)**—Displays the time interval you have set for page refresh.
- **Clear Statistics**—Clears the statistics for the interface selected from the table.

Using the CLI:

- 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 In the J-Web interface the charts displayed are:

- Bar charts—Display the input and output error counters.
- Pie charts—Display the number of broadcast, unicast, and multicast packet counters.

For details about output from the CLI commands, see **show interfaces ge-** (Gigabit Ethernet) or **show interfaces xe-** (10-Gigabit Ethernet).

Related Documentation

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

Verifying the Status of a LAG Interface

Purpose Verify that a LAG (ae0) has been created on the switch.

Action Enter the following command:

```
user@switch> 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 Aggregated Ethernet Links (CLI Procedure) on page 121
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 122
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33

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.

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-0/1/0
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 ouput shows that LACP has been configured with one side as active and the other as passive. When LACP is enabled, at least one side must be set as active 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 interfaces ae*x* statistics** command to display LACP BPDU exchange information.

show interfaces ae0 statistics

```
Physical interface: ae0, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 30
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
    Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet,
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no PDUs are being exchanged (when there is no other traffic flowing on the link).

Related Documentation

- *Configuring Aggregated Ethernet LACP*
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- *Verifying the Status of a LAG Interface*
- Verifying the Status of a LAG Interface on page 291

Verifying That Layer 3 Subinterfaces Are Working

Purpose After configuring Layer 3 subinterfaces, verify they are set up properly and transmitting data.

Action

1. Use the **show interfaces** command to determine whether you successfully created the subinterfaces and the links are up:

```
user@switch> show interfaces interface-name terse
Interface               Admin Link Proto    Local                 Remote
ge-0/0/0                up    up
ge-0/0/0.0              up    up   inet     1.1.1.1/24
ge-0/0/0.1              up    up   inet     2.1.1.1/24
ge-0/0/0.2              up    up   inet     3.1.1.1/24
ge-0/0/0.3              up    up   inet     4.1.1.1/24
ge-0/0/0.4              up    up   inet     5.1.1.1/24
ge-0/0/0.32767          up    up
```

2. Use the **ping** command from a device on one subnet to an address on another subnet to determine whether packets were transmitted correctly on the subinterface VLANs:

```
user@switch> ping ip-address
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
```

Meaning The output confirms that the subinterfaces are created and the links are up.

Related Documentation

- Configuring a Layer 3 Subinterface (CLI Procedure) on page 135
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 44

Verifying Unicast RPF Status

Purpose Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

Action Use one of the **show interfaces** ***interface-name*** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The following example displays output from the **show interfaces ge- extensive** command.

```
user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
  Interface index: 139, SNMP ifIndex: 58, Generation: 140
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  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: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
Last flapped   : Never
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, 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  : 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
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: 0, CAM source filters: 0
Autonegotiation information:
```



```
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   IPv6 transit 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
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
      Protocol inet, Generation: 144, Route table: 0
    Flags: uRPF
    Addresses, Flags: Is-Preferred Is-Primary
```

Meaning The **show interfaces ge-1/0/10 extensive** command (and the **show interfaces ge-1/0/10 detail** command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On EX3200, EX4200, and EX4300 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs) when you enable unicast RPF on a single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on EX3200 and EX4200 switches.

Related Documentation

- **show interfaces xe- on page 375**
- Example: Configuring Unicast RPF on an EX Series Switch on page 51
- Configuring Unicast RPF (CLI Procedure) on page 135
- Disabling Unicast RPF (CLI Procedure) on page 137

- Troubleshooting Unicast RPF on page 435

Verifying IP Directed Broadcast Status

Purpose Verify that IP directed broadcast is enabled and is working on the subnet.

Action Use the **show vlans extensive** command to verify that IP directed broadcast is enabled and working on the subnet as shown in "Example: Configuring IP Directed Broadcast on an EX Series Switch" on page 56.

Related Documentation

- Configuring IP Directed Broadcast (CLI Procedure) on page 137
- *Configuring IP Directed Broadcast (CLI Procedure)*
- Example: Configuring IP Directed Broadcast on an EX Series Switch on page 56

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) on page 142

Verifying That EEE Is Saving Energy on Configured Ports

Purpose Verify that enabling EEE saves energy on Base-T Copper Ethernet ports.

Action You can see the amount of energy saved by EEE on an EX Series switch using the **show chassis power-budget-statistics** command.

1. View the power budget of an EX Series switch before enabling EEE.

   - On an EX6210 switch:

```
user@switch>show chassis power-budget-statistics
     PSU  2     (EX6200-PWR-AC2500)                :    2500 W   Online
            PSU  3                 )                :       0 W  Offline
            Total Power supplied by all Online PSUs  :    2500 W
            Power Redundancy Configuration          :     N+1
            Power Reserved for the Chassis          :     500 W
    Fan Tray Statistics                          Base power   Power Used
    FTC  0                                      :     300 W        nan W
    FPC Statistics                               Base power   Power Used   PoE
 power   Priority
    FPC  3   (EX6200-48T)                       :     150 W      61.54 W
    0 W       9
    FPC  4   (EX6200-SRE64-4XS)                 :     100 W      48.25 W
    0 W       0
    FPC  5   (EX6200-SRE64-4XS)                 :     100 W      48.00 W
    0 W       0
    FPC  7   (EX6200-48T)                       :     150 W      63.11 W
    0 W       9
    FPC  8   (EX6200-48T)                       :     150 W      12.17 W
    0 W       9

    Total (non-PoE) Power allocated             :     950 W
            Total Power allocated for PoE            :       0 W
            Power Available (Redundant case)         :       0 W
            Total Power Available                    :    1550 W
```

   - On an EX4300 switch:

```
user@switch>show chassis power-budget-statistics fpc 1
     PSU  1    (JPSU-1100-AC-AFO-A)         :    1100 W   Online
            Power redundancy configuration          :    N+0
            Total power supplied by all online PSUs  :    1100 W
            Base power reserved                      :    175 W
            Non-PoE power being consumed             :    95 W
            Total Power allocated for PoE            :    925 W
            Total PoE power consumed                 :    0 W
            Total PoE power remaining                :    925 W
```

2. Enable EEE on Base-T Copper Ethernet ports and save the configuration.
3. View the power budget of the switch after enabling EEE.

   - On an EX6210 switch:

```
user@switch> show chassis power-budget-statistics
     PSU  2     (EX6200-PWR-AC2500)                :    2500 W   Online
            PSU  3                 )                :       0 W  Offline
            Total Power supplied by all Online PSUs  :    2500 W
            Power Redundancy Configuration          :     N+1
            Power Reserved for the Chassis          :     500 W
    Fan Tray Statistics                          Base power   Power Used
    FTC  0                                      :     300 W        nan W
    FPC Statistics                               Base power   Power Used   PoE
```

```
power   Priority
   FPC  3   (EX6200-48T)                 :     150 W      50.36 W
   0 W       9
   FPC  4   (EX6200-SRE64-4XS)           :     100 W      48.60 W
   0 W       0
   FPC  5   (EX6200-SRE64-4XS)           :     100 W      48.09 W
   0 W       0
   FPC  7   (EX6200-48T)                 :     150 W      51.38 W
   0 W       9
   FPC  8   (EX6200-48T)                 :     150 W      12.17 W
   0 W       9

   Total (non-PoE) Power allocated       :     950 W
           Total Power allocated for PoE         :       0 W
           Power Available (Redundant case)      :       0 W
           Total Power Available                 :    1550 W
```

- On an EX4300 switch:

```
user@switch> show chassis power-budget-statistics fpc 1
     PSU  1    (JPSU-1100-AC-AFO-A)         :    1100 W   Online
            Power redundancy configuration        :    N+0
            Total power supplied by all online PSUs  :    1100 W
            Base power reserved                   :    175 W
            Non-PoE power being consumed          :    86 W
            Total Power allocated for PoE         :    925 W
            Total PoE power consumed              :    0 W
            Total PoE power remaining             :    925 W
```

Meaning On an EX6210 switch, the **Power Used** field in the output shows the actual power being consumed by the line card or SRE module, including PoE power. If you compare the values in the **Power Used** field before and after enabling EEE for FPC 3 and FPC 7, you will notice that power is saved when EEE is enabled.

NOTE: The Power Used field is displayed in the output only for EX6210 switches.

On an EX4300 switch, if you compare the values in the **Non-PoE power being consumed** field before and after enabling EEE, you will notice that power is saved when EEE is enabled.

Related Documentation

- Configuring Energy Efficient Ethernet on Interfaces (CLI Procedure) on page 145
- Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 27

CHAPTER 6

Operational Commands

Common Output Fields Description

This chapter explains the content of the output fields, which appear in the output of most **show interfaces** commands.

Damping Field

For the physical interface, the Damping field shows the setting of the following damping parameters:

- **half-life**—Decay half-life. The number of seconds after which the accumulated interface penalty counter is reduced by half if the interface remains stable.
- **max-suppress**—Maximum hold-down time. The maximum number of seconds that an interface can be suppressed no matter how unstable the interface has been.
- **reuse**—Reuse threshold. When the accumulated interface penalty counter falls below this number, the interface is no longer suppressed.

- **suppress**—Cutoff (suppression) threshold. When the accumulated interface penalty counter exceeds this number, the interface is suppressed.
- **enable**—Enable damping on a per-interface basis. If damping is enabled on an interface, it is suppressed during interface flaps that match the configuration settings.

Destination Class Field

For the logical interface, the **Destination class** field provides the names of destination class usage (DCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

```
                                      Packets                         Bytes
Destination class              (packet-per-second)            (bits-per-second)

            gold                      1928095                  161959980
                               (             889)             (    597762)
           bronze                           0                          0
                               (               0)             (          0)
           silver                           0                          0
                               (               0)             (          0)
```

Enabled Field

For the physical interface, the **Enabled** field provides information about the state of the interface, displaying one or more of the following values:

- **Administratively down, Physical link is Down**—The interface is turned off, and the physical link is inoperable and cannot pass packets even when it is enabled.To change the interface state to **Enabled**, use the following command:

  user@host# set interfaces *interface* enable

  Manually verify the connections to bring the physical link up.
- **Administratively down, Physical link is Up**—The interface is turned off, but the physical link is operational and can pass packets when it is enabled.To change the interface state to **Enabled**, use the following command:

  user@host# set interfaces *interface* enable
- **Enabled, Physical link is Down**—The interface is turned on, but the physical link is inoperable and cannot pass packets. Manually verify the connections to bring the physical link up.
- **Enabled, Physical link is Up**—The interface is turned on, and the physical link is operational and can pass packets.

Filters Field

For the logical interface, the **Filters** field provides the name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is **Filters: Input:** *filter-name* and **Filters: Output:** *filter-name*. For example:

Filters: Input: sample-all

Filters: Output: cp-ftp

Flags Fields

The following sections provide information about flags that are specific to interfaces:

Addresses, Flags Field

The **Addresses, Flags** field provides information about the addresses configured for the protocol family on the logical interface and displays one or more of the following values:

- **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status
- **Is-Default**—The default address of the router used as the source address by SNMP, ping, traceroute, and other network utilities.
- **Is-Preferred**—The default local address for packets originating from the local router and sent to destinations on the subnet.
- **Is-Primary**—The default local address for broadcast and multicast packets originated locally and sent out the interface.
- **Preferred**—This address is a candidate to become the preferred address.
- **Primary**—This address is a candidate to become the primary address.

Device Flags Field

The **Device flags** field provides information about the physical device and displays one or more of the following values:

- **Down**—Device has been administratively disabled.
- **Hear-Own-Xmit**—Device receives its own transmissions.
- **Link-Layer-Down**—The link-layer protocol has failed to connect with the remote endpoint.
- **Loopback**—Device is in physical loopback.
- **Loop-Detected**—The link layer has received frames that it sent, thereby detecting a physical loopback.
- **No-Carrier**—On media that support carrier recognition, no carrier is currently detected.
- **No-Multicast**—Device does not support multicast traffic.
- **Present**—Device is physically present and recognized.

- **Promiscuous**—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media.
- **Quench**—Transmission on the device is quenched because the output buffer is overflowing
- **Recv-All-Multicasts**—Device is in multicast promiscuous mode and therefore provides no multicast filtering.
- **Running**—Device is active and enabled.

Family Flags Field

The **Family flags** field provides information about the protocol family on the logical interface and displays one or more of the following values:

- **DCU**—Destination class usage is enabled.
- **Dest-route-down**—The software detected that the link is down and has stopped forwarding the link's interface routes.
- **Down**—Protocol is inactive.
- **Is-Primary**—Interface is the primary one for the protocol.
- **Mac-Validate-Loose**—Interface is enabled with loose MAC address validation.
- **Mac-Validate-Strict**—Interface is enabled with strict MAC address validation.
- **Maximum labels**—Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.
- **MTU-Protocol-Adjusted**—The effective MTU is not the configured value in the software.
- **No-Redirects**—Protocol redirects are disabled.
- **Primary**—Interface can be considered for selection as the primary family address.
- **Protocol-Down**—Protocol failed to negotiate correctly.
- **SCU-in**—Interface is configured for source class usage input.
- **SCU-out**—Interface is configured for source class usage output.
- **send-bcast-packet-to-re**—Interface is configured to forward IPv4 broadcast packets to the Routing Engine.
- **targeted-broadcast**—Interface is configured to forward IPv4 broadcast packets to the LAN interface and the Routing Engine.
- **Unnumbered**—Protocol family is configured for unnumbered Ethernet. An unnumbered Ethernet interface borrows an IPv4 address from another interface, which is referred to as the donor interface.
- **Up**–Protocol is configured and operational.
- **uRPF**—Unicast Reverse Path Forwarding is enabled.

Interface Flags Field

The **Interface flags** field provides information about the physical interface and displays one or more of the following values:

- **Admin-Test**—Interface is in test mode and some sanity checking, such as loop detection, is disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Hardware-Down**—Interface is nonfunctional or incorrectly connected.
- **Link-Layer-Down**—Interface keepalives have indicated that the link is incomplete.
- **No-Multicast**—Interface does not support multicast traffic.
- **No-receive No-transmit**—Passive monitor mode is configured on the interface.
- **OAM-On-SVLAN**—(MX Series routers with MPC/MIC interfaces only) Interface is configured to propagate the Ethernet OAM state of a static, single-tagged service VLAN (S-VLAN) on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interface to a dynamic or static double-tagged customer VLAN (C-VLAN) that has the same S-VLAN (outer) tag as the S-VLAN.
- **Point-To-Point**—Interface is point-to-point.
- **Pop all MPLS labels from packets of depth**—MPLS labels are removed as packets arrive on an interface that has the **pop-all-labels** statement configured. The depth value can be one of the following:
  - **1**—Takes effect for incoming packets with one label only.
  - **2**—Takes effect for incoming packets with two labels only.
  - **[1 2]**—Takes effect for incoming packets with either one or two labels.
- **Promiscuous**—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses.
- **Recv-All-Multicasts**—Interface is in multicast promiscuous mode and provides no multicast filtering.
- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Link Flags Field

The **Link flags** field provides information about the physical link and displays one or more of the following values:

- **ACFC**—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option.
- **Give-Up**—Link protocol does not continue connection attempts after repeated failures.

- **Loose-LCP**—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational.
- **Loose-LMI**—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational.
- **Loose-NCP**—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational.
- **No-Keepalives**—Link protocol keepalives are disabled.
- **PFC**—Protocol field compression is configured. The PPP session negotiates the PFC option.

Logical Interface Flags Field

The **Logical interface flags** field provides information about the logical interface and displays one or more of the following values:

- **ACFC Encapsulation**—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer).
- **Device-down**—Device has been administratively disabled.
- **Disabled**—Interface is administratively disabled.
- **Down**—A hardware failure has occurred.
- **Clear-DF-Bit**—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit.
- **Hardware-Down**—Interface protocol initialization failed to complete successfully.
- **PFC**—Protocol field compression is enabled for the PPP session.
- **Point-To-Point**—Interface is point-to-point.
- **SNMP-Traps**—SNMP trap notifications are enabled.
- **Up**—Interface is enabled and operational.

Label-Switched Interface Traffic Statistics Field

When you use the **vrf-table-label** statement to configure a VRF routing table, a label-switched interface (LSI) logical interface label is created and mapped to the VRF routing table.

Any routes present in a VRF routing table and configured with the **vrf-table-label** statement are advertised with the LSI logical interface label allocated for the VRF routing table. When packets for this VPN arrive on a core-facing interface, they are treated as if the enclosed IP packet arrived on the LSI interface and are then forwarded and filtered based on the correct table. For more information on the **vrf-table-label** statement, including a list of supported interfaces, see the *Junos VPNs Configuration Guide*.

If you configure the **family mpls** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level and you also configure the **vrf-table-label** statement at the **[edit routing-instances *routing-instance-name*]** hierarchy level, the output for the

show interface ***interface-name*** **extensive** command includes the following output fields about the LSI traffic statistics:

- **Input bytes**—Number of bytes entering the LSI and the current throughput rate in bits per second (bps).
- **Input packets**—Number of packets entering the LSI and the current throughput rate in packets per second (pps).

NOTE: If LSI interfaces are used with VPLS when **no-tunnel-services** is configured or L3VPN when **vrf-table-label** configuration is applied inside the routing-instance, the **Input packets** field associated with the core-facing interfaces may not display the correct value. Only the Input counter is affected because the LSI is used to receive traffic from the remote PEs. Traffic that arrives on an LSI interface might not be counted at both the Traffic Statistics and the Label-switched interface (LSI) traffic statistics levels.

This note applies to the following platforms:

- M Series routers with -E3 FPC model numbers or configured with an Enhanced CFEB (CFEB-E), and M120 routers
- MX Series routers with DPC or ADPC only

The following example shows the LSI traffic statistics that you might see as part of the output of the **show interface** ***interface-name*** **extensive** command:

```
Label-switched interface (LSI) traffic statistics:
  Input bytes:          0          0 bps
  Input packets:         0          0 pps
```

Policer Field

For the logical interface, the **Policer** field provides the policers that are to be evaluated when packets are received or transmitted on the interface. The format is **Policer: Input:** ***type-fpc/picport*****-in-policer, Output:** ***type-fpc/pic/port*****-out-policer**. For example:

Policer: Input: at-1/2/0-in-policer, Output: at-2/4/0-out-policer

Protocol Field

For the logical interface, the **Protocol** field indicates the protocol family or families that are configured on the interface, displaying one or more of the following values:

- **aenet**—Aggregated Ethernet. Displayed on Fast Ethernet interfaces that are part of an aggregated Ethernet bundle.
- **ccc**—Circuit cross-connect (CCC). Configured on the logical interface of CCC physical interfaces.

- **inet**—IP version 4 (IPv4). Configured on the logical interface for IPv4 protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—IP version 6 (IPv6). Configured on the logical interface for IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), and BGP.
- **iso**—International Organization for Standardization (ISO). Configured on the logical interface for IS-IS traffic.
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI). Configured on the logical interface for link services bundling.
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end. Configured on the logical interface for multilink bundling.
- **mlppp**—Multilink Point-to-Point Protocol (MLPPP). Configured on the logical interface for multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). Configured on the logical interface for participation in an MPLS path.
- **pppoe**— Point-to-Point Protocol over Ethernet (PPPoE). Configured on Ethernet interfaces enabled to support multiple protocol families.
- **tcc**—Translational cross-connect (TCC). Configured on the logical interface of TCC physical interfaces.
- **tnp**—Trivial Network Protocol (TNP). Used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only.
- **vpls**—Virtual private LAN service (VPLS). Configured on the logical interface on which you configure VPLS.

RPF Failures Field

For the logical interface, the **RPF Failures** field provides information about the amount of incoming traffic (in packets and bytes) that failed a unicast reverse path forwarding (RPF) check on a particular interface. The format is **RPF Failures: Packets: xx,Bytes: yy**. For example:

```
RPF Failures: Packets: 0, Bytes:0
```

Source Class Field

For the logical interface, the **Source class** field provides the names of source class usage (SCU) counters per family and per class for a particular interface. The counters display packets and bytes arriving from designated user-selected prefixes. For example:

```
                       Packets             Bytes
Source class      (packet-per-second)     (bits-per-second)

              gold      1928095         161959980
                  (         889)     (     597762)
```



```
    bronze              0                         0
         (             0)          (             0)
     silver             0                         0
         (             0)          (             0)
```

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 39 on page 308. The keys are not case-sensitive.

Table 39: Output Control Keys for the monitor interface Command

Key	Action
c	Clears (returns to zero) the delta counters since **monitor interface** was started. This does not clear the accumulative counter. To clear the accumulative counter, use the **clear interfaces interval** 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 39: Output Control Keys for the monitor interface Command *(continued)*

Key	Action
n	Displays information about the next interface. The **monitor interface** command displays the physical or logical interfaces in the same order as the **show interfaces terse** 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 40 on page 309. The keys are not case-sensitive.

Table 40: 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 **Current Delta** column. The statistics counters are not cleared.
d	Displays the **Current Delta** 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 **Current Delta** column) in Bps and pps.

Required Privilege Level

trace

List of Sample Output

monitor interface (Physical) on page 311
monitor interface (OTN Interface) on page 312
monitor interface (MX2020 Routers with MPC6E and OTN MICInterface) on page 313
monitor interface (Logical) on page 314
monitor interface (QFX3500 Switch) on page 314
monitor interface traffic on page 315
monitor interface traffic (QFX3500 Switch) on page 315
monitor interface traffic detail (QFX3500 Switch) on page 316

Output Fields

Table 41 on page 310 describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 41: monitor interface Output Fields

Field Name	Field Description	Level of Output
router1	Hostname of the router.	All levels
Seconds	How long the monitor interface command has been running or how long since you last cleared the counters.	All levels
Time	Current time (UTC).	All levels
Delay *x/y/z*	Time difference between when the statistics were displayed and the actual clock time. • *x*—Time taken for the last polling (in milliseconds). • *y*—Minimum time taken across all pollings (in milliseconds). • *z*—Maximum time taken across all pollings (in milliseconds).	All levels
Interface	Short description of the interface, including its name, status, and encapsulation.	All levels
Link	State of the link: **Up**, **Down**, or **Test**.	All levels
Current delta	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
Local Statistics	(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. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	All levels
Remote Statistics	(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. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	All levels

Table 41: monitor interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	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. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	All levels
Description	With the **traffic** option, displays the interface description configured at the **[edit interfaces *interface-name*]** 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: OC48
Traffic statistics:                                            Current Delta
  Input packets:                     6045 (0 pps)                       [11]
  Input bytes:                    6290065 (0 bps)                    [13882]
  Output packets:                   10376 (0 pps)                       [10]
  Output bytes:                  10365540 (0 bps)                     [9418]
Encapsulation statistics:
  Input keepalives:                  1901                                [2]
  Output keepalives:                 1901                                [2]
  NCP state: Opened
  LCP state: Opened
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:                  1                                [0]
  Output errors:                        0                                [0]
  Output drops:                         0                                [0]
  Aged packets:                         0                                [0]
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                             1                                [0]
  LOF count                             1                                [0]
  SEF count                             1                                [0]
  ES-S                                  0                                [0]
  SES-S                                 0                                [0]
SONET statistics:
  BIP-B1                           458871                                [0]
```



```
  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    APSPCC 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                                                      [0]
    BIP                               0
    BBE                               0                                   [0]
    ES                                0                                   [0]
    SES                               0                                   [0]
    UAS                               0                                   [0]
  OTN ODU - Counters                                                      [0]
    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:                                               Current delta
  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:                                             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:                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    Link  Input packets        (pps)     Output packets        (pps)
Description
 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
```

request diagnostics tdr

Syntax request diagnostics tdr (abort | start) interface *interface-name*

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Start a time domain reflectometry (TDR) diagnostic test on the specified interface. This test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 interfaces on line cards.

NOTE: **We recommend running the TDR test when there is no traffic on the interface under test.**

You view the results of the TDR test with the **show diagnostics tdr** command.

Options **abort**—Stop the TDR test currently in progress on the specified interface. No results are reported, and previous results, if any, are cleared.

interface-name—The name of the interface.

start—Start a TDR test on the specified interface.

Required Privilege Level maintenance

Related Documentation

- **show diagnostics tdr on page 319**
- Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) on page 436

List of Sample Output **request diagnostics tdr start interface ge-0/0/19 on page 318**

Output Fields Table 42 on page 318 lists the output fields for the **request diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 42: request diagnostics tdr Output Fields

Field Name	Field Description
Test Status	Information about the status of the TDR test request: • **Admin Down** ***interface-name***—The interface is administratively down. The TDR test cannot run on interfaces that are administratively down. • **Interface** ***interface-name*** **not found**—The interface does not exist. • **Test successfully executed** ***interface-name***—The test has successfully started on the interface. You can view the test results with the **show diagnostics tdr** command. • **VCT not supported on** ***interface-name***—The TDR test is not supported on the interface.

Sample Output

request diagnostics tdr start interface ge-0/0/19

```
user@switch> request diagnostics tdr start interface ge-0/0/19

Interface TDR detail:
Test status                  : Test successfully executed  ge-0/0/19
```

show diagnostics tdr

Syntax

show diagnostics tdr
<interface *interface-name*>

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the results of a time domain reflectometry (TDR) diagnostic test run on an interface. A TDR test characterizes and locates faults on twisted-pair Ethernet cables. For example, it can detect a broken twisted pair and provide the approximate distance to the break. It can also detect polarity swaps, pair swaps, and excessive skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches— RJ-45 interfaces on line cards.

Use the request diagnostics tdr command to request a TDR test on a specified interface. Use the **show diagnostic tdr** command to display the last TDR test results for a specified interface or the last TDR test results for all network interfaces on the switch that support the TDR test.

Options

none—Show summarized last results for all interfaces on the switch that support the TDR test.

interface *interface-name*—(Optional) Show detailed last results for the specified interface or a range of interfaces. Specify a range of interfaces by entering the beginning and ending interface in the range, separated by a dash—for example, **ge-0/0/15-ge-0/0/20**.

Required Privilege Level

view

Related Documentation

- **request diagnostics tdr on page 317**
- Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) on page 436

List of Sample Output

show diagnostics tdr interface ge-0/0/19 (Normal Cable) on page 321
show diagnostics tdr interface ge-2/0/2 (Faulty Cable) on page 322
show diagnostics tdr (All Supported Interfaces) on page 322

Output Fields

Table 43 on page 320 lists the output fields for the **show diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 43: show diagnostics tdr Output Fields

Field Name	Field Description
Interface name or **Interface**	Name of interface for which TDR test results are being reported.
Test status	Status of TDR test: • **Aborted**—Test was terminated by operator before it was complete. • **Failed**—Test was not completed successfully. • **Interface *interface-name* not found**—Specified interface does not exist. • **Not Started**—No TDR test results are available for the interface. • **Passed**—Test completed successfully. The cable, however, might still have a fault—see the **Cable status** field for information on the cable. • **Started**—Test is currently running and not yet complete. • **VCT not supported on *interface-name***—TDR test is not supported on the interface.
Link status	Operating status of link: **UP** or **Down**.
MDI pair	Twisted pair for which test results are being reported, identified by pin numbers. (Displayed only when the **interface** option is used.)

Cable status	When detailed information is displayed, status for a twisted pair:
	
	• **Failed**—TDR test failed on the cable pair.
	• **Impedance Mismatch**—Impedance on the twisted pair is not correct. Possible reasons for an impedance mismatch include:
	  ◦ The twisted pair is not connected properly.
	  ◦ The twisted pair is damaged.
	  ◦ The connector is faulty.
	• **Normal**—No cable fault detected for the twisted pair.
	• **Open**—Lack of continuity between the pins at each end of the twisted-pair.
	• **Short on Pair-*n***—A short-circuit was detected on the twisted pair.
	
	When summary information for all interfaces is displayed, status for the cable as a whole:
	
	• **Fault**—A fault was detected on one or more of the twisted-pairs.
	• **OK**—No fault was detected on any of the twisted pairs.

Distance fault or **Max distance fault**	Distance to the fault in whole meters. If there is no fault, this value is 0. When summary information for all interfaces is displayed, this value is the distance to the most distant fault if there is more than one twisted pair with a fault.

Table 43: show diagnostics tdr Output Fields *(continued)*

Field Name	Field Description
Polarity swap	Indicates the polarity status of the twisted pair: • **Normal**—Polarity is normal. Each conductor in the twisted pair has been connected the same pins at the both ends of the connection. For example, a conductor connected to pin 1 at the near end of the connection is connected to pin 1 at the far end. • **Reversed**—Polarity has been reversed. For the twisted pair, the conductors have switched which pins they are connected to at the near and far ends of the connection. For example, the conductor connected to pin 1 at the near end is connected to pin 2 at the far end. (Not available on EX8200 switches.) (Displayed only when the **interface** option is used)
Skew time	Difference in nanoseconds between the propagation delay on this twisted pair and the twisted pair with the shortest propagation delay. (Not available on EX8200 switches.) (Displayed only when the **interface** option is used.)
Channel Pair	Number of the 10/100BASE-T transmit/receive pair being reported on.
Pair Swap	Indicates whether or not the twisted pairs are swapped: • **MDI**—The pairs are not swapped (straight-through cable). • **MDIX**—The pairs are swapped (cross-over cable). (Displayed only when the **interface** option is used.)
Downshift	Indicates whether the connection speed is being downshifted: • **No Downshift**—No downshifting of connection speed. • **Downshift occurs**—Connection speed is downshifted to 10 or 100 Mbs. This occurs if the cable is a two-pair cable rather than the four-pair cable required by Gigabit Ethernet. (Displayed only when the **interface** option is used.)

Sample Output

show diagnostics tdr interface ge-0/0/19 (Normal Cable)

```
user@switch> show diagnostics tdr interface ge-0/0/19
Interface TDR detail:
Interface name                  : ge-0/0/19
Test status                     : Passed
Link status                     :  UP
MDI pair                        : 1-2
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : Normal
  Skew time                     : 0 ns
MDI pair                        : 3-6
```

```
  Cable status                : Normal
  Distance fault              : 0 Meters
  Polartiy swap               : Normal
  Skew time                   : 8 ns
MDI pair                      : 4-5
  Cable status                : Normal
  Distance fault              : 0 Meters
  Polartiy swap               : Normal
  Skew time                   : 8 ns
MDI pair                      : 7-8
  Cable status                : Normal
  Distance fault              : 0 Meters
  Polartiy swap               : Normal
  Skew time                   : 8 ns
Channel pair                  : 1
  Pair swap                   : MDI
Channel pair                  : 2
  Pair swap                   : MDI
Downshift                     : No Downshift
```

show diagnostics tdr interface ge-2/0/2 (Faulty Cable)

```
user@switch> show diagnostics tdr interface ge-2/0/2
Interface TDR detail:
Interface name                : ge-2/0/2
Test status                   : Passed
Link status                   : Down
MDI Pair                      : 1-2
   Cable status               : 1-2
   Distance fault             : 2 Meters
   Polartiy swap              : N/A
   Skew time                  : N/A
MDI Pair                      : 3-6
   Cable status               : Impedance Mismatch
   Distance fault             : 3 Meters
   Polartiy swap              : N/A
   Skew time                  : N/A
MDI Pair                      : 4-5
   Cable status               : Impedance Mismatch
   Distance fault             : 3 Meters
   Polartiy swap              : N/A
   Skew time                  : N/A
MDI Pair                      : 7-8
   Cable status               : Short on Pair-2
   Distance fault             : 3 Meters
   Polartiy swap              : N/A
   Skew time                  : N/A
Channel pair                  : 1
   Pair swap                  : N/A
Channel pair                  : 2
   Pair swap                  : N/A
Downshift                     : N/A
```

show diagnostics tdr (All Supported Interfaces)

```
user@switch> show diagnostics tdr
Interface   Test status    Link status  Cable status  Max distance fault
 ge-0/0/0   Not Started    N/A          N/A           N/A
 ge-0/0/1   Not Started    N/A          N/A           N/A
 ge-0/0/2       Started    N/A          N/A           N/A
 ge-0/0/3       Started    N/A          N/A           N/A
```



```
 ge-0/0/4          Passed     UP          OK           0
 ge-0/0/5          Passed     UP          Fault      173
 ge-0/0/6          Passed     UP          OK           0
 ge-0/0/7          Passed     UP          OK           0
 ge-0/0/8          Passed     UP          OK           0
 ge-0/0/9          Passed     UP          OK           0
ge-0/0/10          Passed     UP          OK           0
ge-0/0/11          Passed     UP          OK           0
ge-0/0/12          Passed     UP          OK           0
ge-0/0/13          Passed     UP          OK           0
ge-0/0/14          Passed     UP          OK           0
ge-0/0/15          Passed     UP          OK           0
ge-0/0/16          Passed     UP          OK           0
ge-0/0/17          Passed     UP          OK           0
ge-0/0/18          Passed     UP          OK           0
ge-0/0/19          Passed     UP          OK           0
ge-0/0/20          Passed     Down        Fault        0
ge-0/0/21          Passed     Down        Fault        5
ge-0/0/22          Passed     UP          OK           0
ge-0/0/23          Passed     UP          OK           0
```

show interfaces (Aggregated Ethernet)

Syntax

show interfaces ae *number*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

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 14.1 for PTX Series Packet Transport Routers.

Description

(M Series, T Series, MX Series, and PTX Series routers and EX Series switches) Display status information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.

Options

ae *number*—Display standard information about the specified aggregated Fast Ethernet or Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information.

snmp-index *snmp-index*—(Optional) Display information about the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces (Aggregated Ethernet) on page 330
show interfaces brief (Aggregated Ethernet) on page 330
show interfaces detail (Aggregated Ethernet) on page 330
show interfaces extensive (Aggregated Ethernet) on page 331
show interfaces extensive (Aggregated Ethernet with VLAN Stacking) on page 332

Output Fields

Table 44 on page 324 lists the output fields for the **show interfaces** (Aggregated Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 44: show interfaces (Aggregated Ethernet) Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface and state of the interface.	All levels
Enabled	State of the physical interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 299.	All levels

Table 44: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interface index	Index number of the physical interface, which reflects its initialization sequence.	All levels
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Minimum links needed	Number of child links that must be operational for the aggregate interface to be operational.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under "Common Output Fields Description" on page 299.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interfaces Flags" section under "Common Output Fields Description" on page 299.	All levels
Current address	Configured MAC address.	**detail extensive**
Hardware address	Hardware MAC address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up or from up to down. The format is **Last flapped:** ***year-month-day hours:minutes:seconds timezone*** **(*hours:minutes:seconds* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive**
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 44: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number of and rate at which bytes and packets are received and transmitted on the physical interface. • **Input bytes**—Number of bytes and rate, in bps, at which bytes are received on the interface. • **Output bytes**—Number of bytes and rate, in bps, at which bytes are transmitted on the interface. • **Input packets**—Number of packets and rate, in pps, at which packets are received on the interface. • **Output packets**—Number of packets and rate, in pps, at which packets are transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface: • **Errors**—Sum of incoming frame aborts and frame check sequence (FCS) errors. • **Drops**—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 random early detection (RED) mechanism. • **Framing errors**—Number of packets received with an invalid FCS. • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**—Number of frames received that are larger than the giant threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail extensive**
Output errors	Output errors on the interface: • **Carrier transitions** —Number of times the interface has gone from **down** to **up**. 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), then the cable, the far-end system, or the PIC is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—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. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**

Table 44: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
Queue counters	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail 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 of 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. Possible values are described in the "Logical Interface Flags Field" section under "Common Output Fields Description" on page 299.	All levels
VLAN-Tag	Tag Protocol Identifier (TPID) and VLAN identifier.	All levels
Demux	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • **Source Family Inet** • **Destination Family Inet**	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels

Table 44: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output

Statistics	Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.
	• **Bundle**—Information about input and output bundle rates.
	• **Link**—(**detail** and **extensive** only) Information about specific links in the aggregate, including link state and input and output rates.
	• **Adaptive Statistics**—(**extensive** only) Information about adaptive load balancing counter statistics.
	  ◦ **Adaptive Adjusts**—Number of times traffic flow imbalance was corrected by implementation of adaptive load balancing.
	  ◦ **Adaptive Scans**—Number of times the link utilization on each member link of the AE bundle was scanned by for adaptive load balancing
	  ◦ **Adaptive Tolerance**—Tolerance level, in percentage, for load imbalance on link utilization on each member link of the AE bundle.
	  ◦ **Adaptive Updates**—Number of times traffic flow load have been updated on an AE bundle.
	• **Marker Statistics**—(**detail** and **extensive** only) Information about 802.3ad marker protocol statistics on the specified links.
	  ◦ **Marker Rx**—Number of valid marker protocol data units (PDUs) received on this aggregation port.
	  ◦ **Resp Tx**—Number of marker response PDUs transmitted on this aggregation port.
	  ◦ **Unknown Rx**—Number of frames received that either carry the slow protocols Ethernet type value (43B.4) but contain an unknown PDU, or are addressed to the slow protocols group MAC address (43B.3) but do not carry the slow protocols Ethernet type.
	  ◦ **Illegal Rx**—Number of frames received that carry the slow protocols Ethernet type value (43B.4) but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).
		detail extensive none

LACP info	Link Aggregation Control Protocol (LACP) information for each aggregated interface.
	• **Role** can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **System priority**—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer.
	• **System identifier**—Actor or partner system ID, encoded as a MAC address.
	• **Port priority**—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer.
	• **Port number**—Port number assigned to the port by the actor or partner, encoded as an unsigned integer.
	• **Port key**—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer.	

Table 44: show interfaces (Aggregated Ethernet) Output Fields *(continued)*

Field Name	Field Description	Level of Output
LACP Statistics	LACP statistics for each aggregated interface. • **LACP Rx**—LACP received counter that increments for each normal hello. • **LACP Tx**—Number of LACP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received. NOTE: For **LACP Rx** and **LACP Tx**, Packet count is updated only on snmp timer expiry (30 secs).	
protocol-family	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under "Common Output Fields Description" on page 299.	**brief**
Protocol	Protocol family configured on the logical interface. Possible values are described in the "Protocol Field" section under "Common Output Fields Description" on page 299.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about protocol family flags. Possible values are described in the "Family Flags Field" section under "Common Output Fields Description" on page 299.	**detail extensive** none
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about address flags. Possible values are described in the "Addresses Flags" section under "Common Output Fields Description" on page 299.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Aggregated Ethernet)

```
user@host> show interfaces ae0
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 60)
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255
```

show interfaces brief (Aggregated Ethernet)

```
user@host> show interfaces ae0 brief
Physical interface: ae0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384

  Logical interface ae0.0
    Flags: SNMP-Traps 16384 Encapsulation: ENET2
    inet  10.100.1.2/24
```

show interfaces detail (Aggregated Ethernet)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59, Generation: 36
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
  Last flapped   : Never
  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
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                7375                 7375                    0
```

```
  1 expedited-fo                0                    0                    0

  2 assured-forw                0                    0                    0

  3 network-cont             2268                 2268                    0


Logical interface ae0.0 (Index 72) (SNMP ifIndex 60) (Generation 18)
  Flags: SNMP-Traps 16384 Encapsulation: ENET2
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Link:
    fe-0/1/0.0
      Input :             0          0             0            0
      Output:             0          0             0            0
    fe-0/1/2.0
      Input :             0          0             0            0
      Output:             0          0             0            0
    fe-0/1/3.0
      Input :             0          0             0            0
      Output:             0          0             0            0
  Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
    fe-0/1/0.0                 0           0            0            0
    fe-0/1/2.0                 0           0            0            0
    fe-0/1/3.0                 0           0            0            0
  Protocol inet, MTU: 1500, Generation: 37, Route table: 0
    Flags: Is-Primary, Mac-Validate-Strict
    Mac-Validate Failures: Packets: 0, Bytes: 0
      Destination: 10.100.1/24, Local: 10.100.1.2, Broadcast: 10.100.1.255,
      Generation: 49
```

show interfaces extensive (Aggregated Ethernet)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 59, Generation: 36
  Link-level type: Ethernet, MTU: 1514, Speed: 300mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Current address: 00:05:85:8b:bf:f0, Hardware address: 00:05:85:8b:bf:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   60                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    1                    0 pps
   Output packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort            7375                 7375                    0

    1 expedited-fo              0                    0                    0
```

```
    2 assured-forw                   0                    0                    0

    3 network-cont                2268                 2268                    0


Logical interface ae0.0 (Index 73) (SNMP ifIndex 563) (Generation 176)
    Flags: Up SNMP-Traps 0x4000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Adaptive Statistics:
        Adaptive Adjusts:          0
        Adaptive Scans  :          0
        Adaptive Updates:          0
    Link:
      fe-1/0/3.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    LACP info:        Role     System             System      Port     Port  Port

                            priority         identifier  priority  number   key

      fe-1/0/3.0     Actor         127  00:24:dc:85:af:f0       127       2     1

      fe-1/0/3.0   Partner         127  00:23:9c:c3:1f:f0       127       1     1

    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      fe-1/0/3.0              3188        3186            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      fe-1/0/3.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 224, Route table: 0
      Flags: Sendbcast-pkt-to-re
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.40.1.0/30, Local: 10.40.1.1, Broadcast: 10.40.1.3,
Generation: 187
    Protocol multiservice, MTU: Unlimited, Generation: 225, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (Aggregated Ethernet with VLAN Stacking)

```
user@host> show interfaces ae0 extensive
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 155, SNMP ifIndex: 48, Generation: 186
  Link-level type: 52, MTU: 1518, Speed: 2000mbps, Loopback: Disabled, Source
filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:12:1e:19:3f:f0, Hardware address: 00:12:1e:19:3f:f0
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2406875                40152 bps
   Output bytes  :              1124470                22056 bps
   Input  packets:                 5307                    5 pps
   Output packets:                13295                   21 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, Giants: 0, Policed discards:
0, Resource errors: 0
 Output errors:
   Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
 Ingress queues: 4 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0 best-effort                    0               859777                    0

   1 expedited-fo                   0                    0                    0

   2 assured-forw                   0                    0                    0

   3 network-cont                   0                    0                    0

 Egress queues: 4 supported, 4 in use
 Queue counters:       Queued packets  Transmitted packets      Dropped packets

   0 best-effort                    0              1897615                    0

   1 expedited-fo                   0                    0                    0

   2 assured-forw                   0                    0                    0

   3 network-cont                   0               662505                    0


 Logical interface ae0.451 (Index 69) (SNMP ifIndex 167) (Generation 601)
   Flags: SNMP-Traps VLAN-Tag [ 0x8100.451 ]  Encapsulation: VLAN-VPLS
   Statistics        Packets        pps         Bytes          bps
   Bundle:
       Input :           289          0         25685          376
       Output:          1698          4        130375         3096
   Link:
     ge-1/2/0.451
       Input :           289          0         25685          376
       Output:             0          0             0            0
     ge-1/2/1.451
       Input :             0          0             0            0
       Output:          1698          4        130375         3096
   Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
     ge-1/2/0.451              0           0            0            0
     ge-1/2/1.451              0           0            0            0
   Protocol vpls, MTU: 1518, Generation: 849, Route table: 3
     Flags: Is-Primary

 Logical interface ae0.452 (Index 70) (SNMP ifIndex 170) (Generation 602)
   Flags: SNMP-Traps VLAN-Tag [ 0x8100.452 ]  Encapsulation: VLAN-VPLS
   Statistics        Packets        pps         Bytes          bps
   Bundle:
       Input :           293          1         26003         1072
       Output:          1694          3        130057         2400
   Link:
     ge-1/2/0.452
```

```
        Input :          293          1        26003         1072
        Output:         1694          3       130057         2400
      ge-1/2/1.452
        Input :            0          0            0            0
        Output:            0          0            0            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-1/2/0.452               0           0            0            0
      ge-1/2/1.452               0           0            0            0
    Protocol vpls, MTU: 1518, Generation: 850, Route table: 3
      Flags: None
...
```


show interfaces (GRE)

Syntax

show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.1 for EX Series switches.
Command introduced in Junos OS Release 13.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display status information about the specified generic routing encapsulation (GRE) interface.

Options

interface-type—On M Series and T Series routers and EX Series switches, the interface type is **gr-*fpc/pic/port***.

brief | detail | extensive | terse—(Optional) Display the specified output level of interface information.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

NOTE: 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 *Junos OS MPLS Applications Library for Routing Devices* and the *Junos OS, Release 14.1*.

Required Privilege Level

view

List of Sample Output

show interfaces (GRE) on page 339
show interfaces brief (GRE) on page 339
show interfaces detail (GRE) on page 339
show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 340
show interfaces extensive (GRE) on page 341

Output Fields

Table 45 on page 336 lists the output fields for the **show interfaces** (GRE) command. Output fields are listed in the approximate order in which they appear.

Table 45: GRE show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under “Common Output Fields Description” on page 299.	All levels
Interface index	Physical interface's index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Device Flags	Information about the physical device. Possible values are described in the “Device Flags” section under “Common Output Fields Description” on page 299.	All levels
Interface Flags	Information about the interface. Possible values are described in the “Interface Flags” section under “Common Output Fields Description” on page 299.	All levels
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none

Table 45: 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	Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under “Common Output Fields Description” on page 299. describe general information about the logical interface. GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field: • **Reassemble-Pkts**—If the **Flags** field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation.	All levels
IP-Header	IP header of the logical interface. If the **tunnel key** statement is configured, this information is included in the **IP Header** entry. GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field: • **df**—If the **IP-Header** field includes this string immediately following the 16 bits of identification information (that is, if **:df:** displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Copy-tos-to-outer-ip-header	Status of type of service (ToS) bits in the GRE packet header: • **On**—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel. • **Off**—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header. NOTE: EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always **Off** in switch output.	**detail extensive**
Gre keepalives configured	Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel. NOTE: EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always **Off** in switch output.	**detail extensive**
Gre keepalives adjacency state	Status of the other end of the GRE tunnel: **Up** or **Down**. 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 45: GRE show interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	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. • **Input rate**—Rate of bits and packets received on the interface. • **Output rate**—Rate of bits and packets transmitted on the interface.	**detail extensive**
Local statistics	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.	**detail extensive**
Transit statistics	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.	**detail extensive** none
Protocol	Protocol family configured on the logical interface, such as **iso**, **inet6**, or **mpls**.	**detail extensive none**
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
MTU	MTU size on the logical interface.	**detail extensive none**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under “Common Output Fields Description” on page 299.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under “Common Output Fields Description” on page 299.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

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 180.20.30.2:180.20.20.3:47: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 339 and show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 340.

show interfaces diagnostics optics

Syntax

show interfaces diagnostics optics *interface-name*

Release Information

Command introduced in Junos OS Release 10.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

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 *Pluggable Transceivers Supported on EX Series Switches*.

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.

Options

interface-name—Name of the interface associated with the port in which the transceiver is installed: ge-*fpc/pic/port*, xe-*fpc/pic/port*, or et-*fpc/pic/port*.

Required Privilege Level

view

Related Documentation

- Monitoring Interface Status and Traffic on page 289
- *Monitoring Interface Status and Traffic*
- *Installing a Transceiver in an EX Series Switch*
- *Installing a Transceiver in a QFX Series Device*
- *Removing a Transceiver from a Switch*
- *Removing a Transceiver from a QFX Series Device*
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

Output Fields

Table 46 on page 342 lists the output fields for the **show interfaces diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

Table 46: show interfaces diagnostics optics Output Fields

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

Table 46: show interfaces diagnostics optics Output Fields *(continued)*

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

Table 46: show interfaces diagnostics optics Output Fields *(continued)*

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

Table 46: show interfaces diagnostics optics Output Fields *(continued)*

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

Table 46: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Module power down alarm (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Displays whether the module power down alarm is **On** or **Off**. When the output is **On**, module is in a limited power mode, low for normal operation.
Tx data not ready alarm (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 **On** or **Off**.
Tx not ready alarm (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 **On** or **Off**.
Tx laser fault alarm (Not available for SFP, SFP+, QSFP+, and CFP transceivers)	Laser fault condition. Displays whether the Tx laser fault alarm is **On** or **Off**.
Tx CDR loss of lock alarm (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 **On** or **Off**.
Rx not ready alarm (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 **On** or **Off**.
Rx loss of signal alarm (Not available for SFP and SFP+ transceivers)	Receive loss of signal alarm. When the output is **On**, indicates insufficient optical input power to the module. Displays whether the Rx loss of signal alarm is **On** or **Off**.
Rx CDR loss of lock alarm (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 **On** or **Off**.
Laser bias current low alarm threshold (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold (Not available for QSFP+ transceivers)	Displays the vendor-specified threshold for the laser output power low alarm.

Table 46: show interfaces diagnostics optics Output Fields *(continued)*

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

Table 46: show interfaces diagnostics optics Output Fields *(continued)*

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

Table 46: show interfaces diagnostics optics Output Fields *(continued)*

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

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
    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 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-

Syntax

show interfaces ge-*fpc/pic/port*
<brief | detail | extensive | terse>
<media>
<statistics>

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display status information about the specified Gigabit Ethernet interface.

NOTE: You must have a transceiver plugged into an SFP or SFP+ port before information about the interface can be displayed.

Options

ge-*fpc/pic/port*—Display standard information about the specified Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

media—(Optional) Display media-specific information about network interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- Monitoring Interface Status and Traffic on page 289
- Troubleshooting Network Interfaces on EX3200 Switches on page 430
- Troubleshooting Network Interfaces on EX4200 Switches on page 432
- Troubleshooting an Aggregated Ethernet Interface on page 429
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

show interfaces ge-0/0/0 on page 363
show interfaces ge-0/0/0 brief on page 363
show interfaces ge-0/0/0 brief (with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces) on page 364
show interfaces ge-0/0/0 detail on page 364
show interfaces ge-0/0/4 extensive on page 365

Output Fields

Table 47 on page 357 lists the output fields for the **show interfaces ge-** command. Output fields are listed in the approximate order in which they appear.

Table 47: show interfaces ge- Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface: **Enabled** or **Disabled**.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Description	Optional user-specified description.	**brief detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. Default is 1514.	All levels
Speed	Speed of the interface: Auto if autonegotiation of speed is enabled; speed in megabits per second if the interface speed is explicitly configured.	All levels
Duplex	Link mode of the interface: Auto if autonegotiation of link mode is enabled; Full-Duplex or Half-Duplex if the link mode is explicitly configured.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
IEEE 802.3az Energy Efficient Ethernet	IEEE 802.3az Energy Efficient Ethernet status: **Enabled** or **Disabled** (appears only for EEE-capable Base-T copper Ethernet interfaces).	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none

Table 47: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago)**.	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface. NOTE: The bandwidth bps counter is not enabled on the switch.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—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. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—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 the Junos OS does not handle. • **L3 incompletes**—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. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—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. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 47: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. 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. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—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. • **Collisions**—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. • **Aged packets**—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. • **FIFO errors**—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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain time, it is promoted to an alarm. Based on the switch configuration, a defect can activate 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 **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—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.	**detail extensive** none

Table 47: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—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). • **FIFO error**—Number of FIFO errors reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—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. • **Fragment frames**—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. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**
Filter Statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem.	**extensive**

Table 47: show interfaces ge- Output Fields *(continued)*

<table>
<tr><th>Field Name</th><th>Field Description</th><th>Level of Output</th></tr>
<tr><td>Autonegotiation information</td><td>Information about link autonegotiation:

Negotiation status:

Complete—The autonegotiation process between the local and remote Ethernet interfaces was successful.
Incomplete—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.
No autonegotiation—Local Ethernet interface has autonegotiation disabled and the link mode and speed are manually configured.

Link partner—Information from the link partner:

Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. If the link mode of the remote device cannot be determined, the value is Unknown.
Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, the types are: Symmetric (link partner supports PAUSE on receive and transmit); Asymmetric (link partner supports PAUSE on transmit); and Symmetric/Asymmetric (link partner supports PAUSE on both receive and transmit or PAUSE only on receive).
Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline.
Link partner speed—Speed of the link partner.

Local resolution—Resolution of the autonegotiation process on the local interface:

Flow control—Type of flow control that is used by the local interface. For Gigabit Ethernet interfaces, the types are: Symmetric (link partner supports PAUSE on receive and transmit); Asymmetric (link partner supports PAUSE on transmit); and Symmetric/Asymmetric (link partner supports PAUSE on both receive and transmit or PAUSE only on receive).
Link mode—Link mode of local interface: either Full-duplex or Half-duplex. Displayed when Negotiation status is Incomplete.
Local link speed—Speed of the local interface. Displayed when Negotiation status is Incomplete.
Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive).

</td><td>extensive</td></tr>
<tr><td>Packet Forwarding Engine configuration</td><td>Information about the configuration of the Packet Forwarding Engine:

Destination slot—FPC slot number:

On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0.
On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch.
On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis.

</td><td>extensive</td></tr>
<tr><td colspan="3">Logical Interface</td></tr>
</table>

Table 47: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface. NOTE: For logical interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic.	**detail extensive**
IPv6 transit statistics	EX Series switches do not support the collection and reporting of IPv6 transit statistics.	**extensive**
Local statistics	Number and rate of bytes and packets destined to and from the switch.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface.	**detail extensive**
Output Filters	Names of any output filters applied to this interface.	**detail extensive**
Flags	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.	**detail extensive**
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flags.	**detail extensive** none

Table 47: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces ge-0/0/0

```
user@switch> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, 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-0/0/0.0 (Index 65) (SNMP ifIndex 22)
    Flags: SNMP-Traps
    Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch
      Flags: None
```

show interfaces ge-0/0/0 brief

```
user@switch> show interfaces ge-0/0/0 brief
Physical interface: ge-0/0/0, 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-0/0/0.0
```



```
    Flags: Device-Down SNMP-Traps Encapsulation: ENET2
    eth-switch
```

show interfaces ge-0/0/0 brief (with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces)

```
user@switch> show interfaces ge-0/0/0 brief
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online,
  IEEE 802.3az Energy Efficient Ethernet: Enabled, NO LPI
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
```

show interfaces ge-0/0/0 detail

```
user@switch> show interfaces ge-0/0/0 detail
Physical interface: ge-0/0/0, 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-0/0/0.0 (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 ge-0/0/4 extensive

```
user@switch> show interfaces ge-0/0/4 extensive
Physical interface: ge-0/0/4, Enabled, Physical link is Up
  Interface index: 165, SNMP ifIndex: 152, Generation: 168
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  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:12:33:65:44, Hardware address: 00:1f:12:33:65:44
  Last flapped   : 2008-09-17 11:02:25 UTC (16:32:54 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  0                    0 bps
   Output bytes  :            2989761                  984 bps
   Input  packets:                  0                    0 pps
   Output packets:              24307                    1 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: 1, 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, 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                24307                    0

  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                             0          2989761
    Total packets                            0            24307
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0            24307
    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
    Code violations                          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: None, Remote fault: Link OK
  Packet Forwarding Engine configuration:
    Destination slot: 0
    Direction : Output
    CoS transmit queue              Bandwidth              Buffer Priority
Limit
                              %            bps     %           usec
    0 best-effort            95      950000000    95             NA      low
none
    7 network-control         5       50000000     5             NA      low
none

  Logical interface ge-0/0/4.0 (Index 82) (SNMP ifIndex 184) (Generation 147)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                  0
     Output bytes  :            4107883
     Input  packets:                  0
     Output packets:              24307
     IPv6 transit statistics:
      Input  bytes  :                 0
      Output bytes  :                 0
      Input  packets:                 0
      Output packets:                 0
    Local statistics:
     Input  bytes  :                  0
     Output bytes  :            4107883
     Input  packets:                  0
     Output packets:              24307
    Transit 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
Protocol eth-switch, Generation: 159, Route table: 0
  Flags: None
  Input Filters: f2,
  Output Filters: f1,,,,
```

show interfaces me0

Syntax

show interfaces me0
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance>
<statistics>

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display status information about the management Ethernet interface.

Options

none—Display standard information about the management Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

routing-instance—(Optional) Display the name of the routing instance.

statistics—(Optional) Display static interface statistics.

Required Privilege Level view

Related Documentation

- *Example: Configuring a Firewall Filter on a Management Interface on an EX Series Switch*
- *Configuring Firewall Filters (CLI Procedure)*

List of Sample Output

show interfaces me0 on page 372
show interfaces me0 brief on page 372
show interfaces me0 detail on page 372
show interfaces me0 extensive on page 373

Output Fields Table 48 on page 368 lists the output fields for the **show interfaces me0** command. Output fields are listed in the approximate order in which they appear.

Table 48: show interfaces me0 Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface: **Enabled** or **Disabled**.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none

Table 48: show interfaces me0 Output Fields *(continued)*

Field Name	Field Description	Level of Output
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Description	Optional user-specified description.	**brief detail extensive**
Type	Information about the type of functional interface.	All levels
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. The default is 1514.	All levels
Clocking	Interface that acts as a clock source. This field is not supported on EX Series switches and the default value is always **Unspecified**.	**detail extensive**
Speed	Speed at which the interface is running.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link type	Information about whether the link is duplex and whether the negotiation is manual or automatic.	**detail extensive** none
Physical info	Information about the device dependent physical interface selector. This field is applied only when a clocking option is specified. This field is not supported on EX Series switches and the default value is always **Unspecified**.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Alternate link address	Information about alternate hardware address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*weeksw:daysdhour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3w:3d 22:58 ago).**	**detail extensive** none
Statistics last cleared	Time when the statistics for the interface was last set to zero. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*weeksw:daysdhour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3w:3d 22:58 ago).**	**detail extensive**

Table 48: show interfaces me0 Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. Following are fields in **Traffic statistics**: • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
IPv6 transit statistics	Number and rate of bytes and IPv6 packets received and transmitted on the physical interface. Following are fields in **IPv6 transit statistics**: • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and frame checksum (FCS) errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid FCS. • **Runts**—Number of frames received that are smaller than the runt threshold. • **Giants**— Number of packets that exceed the size for the medium. For example, if the medium is Ethernet, the **Giant** field shows the count of packets with size greater than 1518 bytes. • **Policed discards**—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 the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**extensive**
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly. It increases 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 increment quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—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. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 48: show interfaces me0 Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	**detail extensive**
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	**detail extensive**
Local statistics	Number and rate of bytes and packets destined to and exiting from the switch.	**extensive**
Protocol	Protocol family.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Routing table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail extensive**
Flags	Information about protocol family flags.	**detail extensive**
Input Filter	Ingress filter name.	**extensive**
Output Filter	Egress filter name.	**extensive**
Addresses	Information about the management interface addresses.	**detail extensive** none
Flags	Information about the address flags.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces me0

```
user@switch> show interfaces me0
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:32 ago)
    Input packets : 1661830
    Output packets: 3200

  Logical interface me0.0 (Index 3) (SNMP ifIndex 34)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 1661830
    Output packets: 3200
    Protocol inet
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.204.32/20, Local: 10.204.33.103,
        Broadcast: 10.204.47.255
    Protocol inet6
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf
```

show interfaces me0 brief

```
user@switch> show interfaces me0 brief
Physical interface: me0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface me0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.204.33.103/20
    inet6 fe80::21f:12ff:fe35:3cbf/64
```

show interfaces me0 detail

```
user@switch> show interfaces me0 detail
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:35:3c:bf, Hardware address: 00:1f:12:35:3c:bf
  Alternate link address: Unspecified
  Last flapped   : 2010-07-31 23:45:50 PDT (5d 00:37 ago)
  Statistics last cleared: Never
```



```
  Traffic statistics:
   Input  bytes  :            366663167
   Output bytes  :               498590
   Input  packets:              1664031
   Output packets:                 3259
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :            366665637
     Output bytes  :               500569
     Input  packets:              1664048
     Output packets:                 3275
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :            366665637
     Output bytes  :               500569
     Input  packets:              1664048
     Output packets:                 3275
    Protocol inet, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
       Destination: 10.204.32/20, Local: 10.204.33.103, Broadcast: 10.204.47.255,
 Generation: 1
    Protocol inet6, Generation: 2, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fe35:3cbf
    Generation: 2
```

show interfaces me0 extensive

```
user@switch> show interfaces me0 extensive
Physical interface: me0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 33, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:38:58:bf, Hardware address: 00:1f:12:38:58:bf
  Alternate link address: Unspecified
  Last flapped   : 2010-08-15 06:27:33 UTC (03:06:22 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :             82310392
   Output bytes  :              1966952
   Input  packets:               110453
   Output packets:                17747
   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, Giants: 0,
  Policed discards: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0,
  Resource errors: 0

Logical interface me0.0 (Index 3) (SNMP ifIndex 34) (Generation 1)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :            82310392
   Output bytes  :             1966952
   Input  packets:              110453
   Output packets:               17747
  Local statistics:
   Input  bytes  :            82310392
   Output bytes  :             1966952
   Input  packets:              110453
   Output packets:               17747
  Protocol inet, Generation: 1, Route table: 0
    Flags: Is-Primary
    Input Filters: mgmt_filter,
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 10.204.96/20, Local: 10.204.96.234,
      Broadcast: 10.204.111.255, Generation: 1
```


show interfaces xe-

Syntax

show interfaces xe-*fpc/pic/port*
<brief | detail | extensive | terse>
<media>
<statistics>

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display status information about the specified 10-Gigabit Ethernet interface.

NOTE: You must have a transceiver plugged into an SFP+ or an XFP port before information about the interface can be displayed.

NOTE: On an EX Series switch, the traffic statistics for a LAG might vary slightly from the cumulative traffic statistics of the member interfaces of the LAG. This difference is more likely to be seen when the traffic is bursty in nature, and because the statistics are not fetched from the LAG and the members in the same instant. For accurate traffic statistics for a LAG, use the aggregated Ethernet counters.

Options

xe-*fpc/pic/port* —Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

media—(Optional) Display media-specific information about network interfaces. For 10-Gigabit Ethernet interfaces, using the **media** option does not provide you with new or additional information. The output is the same as when the **media** option is not used.

statistics—(Optional) Display static interface statistics. For 10-Gigabit Ethernet interfaces, using the **statistics** option does not provide you with new or additional information. The output is the same as when the **statistics** option is not used.

Required Privilege Level view

Related Documentation

- Monitoring Interface Status and Traffic on page 289
- Troubleshooting Network Interfaces on EX3200 Switches on page 430
- Troubleshooting Network Interfaces on EX4200 Switches on page 432
- Troubleshooting an Aggregated Ethernet Interface on page 429
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

show interfaces xe-4/1/0 on page 384
show interfaces xe-0/1/0 brief on page 385
show interfaces xe-4/1/0 detail on page 385
show interfaces xe-6/0/39 extensive on page 386

Output Fields

Table 49 on page 376 lists the output fields for the **show interfaces xe-** command. Output fields are listed in the approximate order in which they appear.

Table 49: show interfaces xe- Output Fields

Field Name	Field Description	Level of Output
Fields for the Terse Output Level Only		
Interface	Name of the physical or logical interface.	**terse**
Admin	Administrative state of the interface.	**terse**
Link	State of the physical link.	**terse**
Proto	Protocol family configured on the logical interface.	**terse**
Local	Local IP address of the logical interface.	**terse**
Remote	Remote IP address of the logical interface.	**terse**
Fields for the Physical Interface		
Physical interface	Name of the physical interface.	**brief** **detail** **extensive** none
Enabled	State of the interface. Can be one of the following: • **Administratively down, Physical link is Down**—The interface is turned off, and the physical link is inoperable and cannot pass packets even when it is enabled. • **Administratively down, Physical link is Up**—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. • **Enabled, Physical link is Down**—The interface is turned on, but the physical link is inoperable and cannot pass packets. • **Enabled, Physical link is Up**—The interface is turned on, and the physical link is operational and can pass packets.	**brief** **detail** **extensive** none
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail** **extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail** **extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail** **extensive**

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Description	User-configured interface description.	**brief** **detail** **extensive** none
Link-level type	Encapsulation being used on the physical interface.	**brief** **detail** **extensive** none
MTU	Maximum transmission unit size on the physical interface.	**brief** **detail** **extensive** none
Speed	Speed at which the interface is running.	**brief** **detail** **extensive** none
Duplex	Duplex mode of the interface.	**brief** **detail** **extensive** none
BPDU Error	Not supported on EX Series switches.	**detail** **extensive** none
MAC-REWRITE Error	Not supported on EX Series switches.	**detail** **extensive** none
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	**brief** **detail** **extensive** none
Source filtering	Source filtering status: **Enabled** or **Disabled**.	**brief** **detail** **extensive** none
Flow control	Flow control status: **Enabled** or **Disabled**.	**brief** **detail** **extensive** none
Device flags	Information about the physical device.	**brief** **detail** **extensive** none

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interface flags	Information about the interface.	**brief** **detail** **extensive** none
Link flags	Information about the link.	**brief** **detail** **extensive** none
CoS queues	Number of CoS queues configured.	**detail** **extensive** none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail** **extensive**
Current address	Configured MAC address.	**detail** **extensive** none
Hardware address	Hardware MAC address.	**detail** **extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is ***year-month-day hour*:*minute*:*second timezone*** **(*weeks*w*days*d *hours*:*minutes*:*seconds* ago)**. For example, **2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail** **extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	none
Output Rate	Output rate in bps and pps.	none
Statistics last cleared	Date, time, and how long ago the statistics for the interface were cleared. The format is ***year-month-day hour*:*minute*:*second timezone*** **(*weeks*w*days*d *hours*:*minutes*:*seconds* ago)**. For example, **2010-05-17 07:51:28 PDT (00:04:33 ago)**.	**detail** **extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface and rate in bits per second. • **Output bytes**—Number of bytes transmitted on the interface and rate in bits per second. • **Input packets**—Number of packets received on the interface and rate in packets per second. • **Output packets**—Number of packets transmitted on the interface and rate in packets per second.	**detail** **extensive**

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	EX Series switches do not support the collection and reporting of IPv6 transit statistics.	**detail extensive**
Input errors	Input errors on the interface: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—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. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—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 the Junos OS does not handle. • **L3 incompletes**—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 **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—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. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. 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. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—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. • **Collisions**—Number of Ethernet collisions. A 10-Gigabit Ethernet interface supports only full-duplex operation, so for 10-Gigabit Ethernet interfaces, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—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. • **FIFO errors**—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. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Ingress queues	Number of CoS ingress queues supported on the specified interface. Displayed only for an interface on a line card with oversubscribed ports.	**detail** **extensive**
Egress queues	Number of CoS egress queues supported on the specified interface.	**detail** **extensive**
PFE Egress queues	Number of Packet Forwarding Engine egress queues shared by the interfaces in a port group. Displayed only for an interface on a line card with oversubscribed ports.	**detail** **extensive**
Queue counters	Statistics for queues: • **Queued packets**—Number of queued packets. This counter is not supported on EX switches and always contains 0. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail** **extensive**

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Active alarms and Active defects	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 **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—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.	**detail** **extensive** none
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—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). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—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. • **Fragment frames**—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. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine:
	
	• **Destination slot**—FPC slot number:
	  ◦ On standalone switches with built-in interfaces, the slot number refers to the switch itself and is always 0.
	  ◦ On Virtual Chassis composed of switches with built-in interfaces, the slot number refers to the member ID of the switch.
	  ◦ On switches with line cards or on Virtual Chassis composed of switches with line cards, the slot number refers to the line card slot number on the switch or Virtual Chassis.
		extensive

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output

CoS Information	Scheduler information for the CoS egress queues on the physical interface:
	• **Direction**—Queue direction, always **Output**.
	• **CoS transmit queue**—Queue number and its associated user-configured forwarding class name.
	• **Bandwidth**—Information about bandwidth allocated to the queue:
	  ◦ **%**—Bandwidth allocated to the queue as a percentage
	  ◦ **bps**—Bandwidth allocated to the queue in bps
	• **Buffer**—Information about buffer space allocated to the queue:
	  ◦ **%**—Buffer space allocated to the queue as a percentage.
	  ◦ **usec**—Buffer space allocated to the queue in microseconds. This value is nonzero only if the buffer size is configured in terms of time.
	• **Priority**—Queue priority: **low** or **high**.
	• **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
		extensive

Fields for MACsec statistics		
Protected Packets	The number of packets sent from the interface that were secured using MACsec when encryption was disabled.	**detail extensive**
Encrypted Packets	The number of packets sent from the interface that were secured and encrypted using MACsec.	**detail extensive**
Protected Bytes	The number of bytes sent from the interface that were secured using MACsec, but not encrypted.	**detail extensive**
Encrypted Bytes	The number of packets sent from the interface that were secured and encrypted using MACsec.	**detail extensive**
Accepted Packets	The number of received packets that have been accepted on the interface. A packet is considered accepted for this counter when it has been received by this interface and it has passed the MACsec integrity check. This counter increments for traffic that is and is not encrypted using MACsec.	**detail extensive**
Validated Bytes	The number of bytes that have been validated by the MACsec integrity check and received on the interface. This counter does not increment when MACsec encryption is disabled.	**detail extensive**
Decrypted Bytes	The number of bytes received on the interface that have been decrypted. An encrypted byte has to be decrypted before it can be received on the receiving interface. The decrypted bytes counter is incremented for received traffic that was encrypted using MACSec.	**detail extensive**
Fields for Logical Interfaces		

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Logical interface	Name of the logical interface.	**brief** **detail** **extensive** none
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**
Description	User-configured description of the interface.	**brief** **detail** **extensive** none
Flags	Information about the logical interface.	**brief** **detail** **extensive** none
Encapsulation	Encapsulation on the logical interface.	**brief** **detail** **extensive** none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface. NOTE: For logical interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic.	**detail** **extensive**
Local statistics	Number and rate of bytes and packets destined to and from the switch.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Protocol	Protocol family.	**detail** **extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail** **extensive**

Table 49: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail** **extensive** none
Input Filters	Names of any input filters applied to this interface.	**detail** **extensive**
Output Filters	Names of any output filters applied to this interface.	**detail** **extensive**
Flags	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.	**detail** **extensive**
Addresses, Flags	Information about the address flags.	**detail** **extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flags.	**detail** **extensive** none
Destination	IP address of the remote side of the connection.	**detail** **extensive** none
Local	IP address of the logical interface.	**detail** **extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail** **extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail** **extensive**

Sample Output

show interfaces xe-4/1/0

```
user@switch show interfaces xe-4/1/0
Physical interface: xe-4/1/0, Enabled, Physical link is Up
  Interface index: 387, SNMP ifIndex: 369
```



```
  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
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
  Last flapped   : 2009-05-12 08:01:04 UTC (00:13:44 ago)
  Input rate     : 36432 bps (3 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch
      Flags: None
```

show interfaces xe-0/1/0 brief

```
user@switch> show interfaces xe-0/1/0 brief
Physical interface: xe-0/1/0, 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-0/1/0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    eth-switch
```

show interfaces xe-4/1/0 detail

```
user@switch> show interfaces xe-4/1/0 detail
Physical interface: xe-4/1/0, Enabled, Physical link is Up
  Interface index: 387, SNMP ifIndex: 369, Generation: 390
  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
  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: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
  Last flapped   : 2009-05-12 08:01:04 UTC (00:13:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              4945644                48576 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                 3258                    4 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 xe-4/1/0.0 (Index 66) (SNMP ifIndex 417) (Generation 158)
    Flags: SNMP-Traps 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, Generation: 174, Route table: 0
      Flags: None
      Input Filters: f1,
      Output Filters: f2,,,,
```

show interfaces xe-6/0/39 extensive

```
user@switch> show interfaces xe-6/0/39 extensive
Physical interface: xe-6/0/39, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 1641, Generation: 316
  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
  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: 00:19:e2:72:f2:88, Hardware address: 00:19:e2:72:f2:88
  Last flapped   : 2010-05-13 14:49:43 PDT (1d 00:14 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :       49625962140160          4391057408 bps
   Output bytes  :       47686985710805          4258984960 bps
   Input  packets:         387702829264             4288139 pps
   Output packets:         372554570944             4159166 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: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 2 supported, 2 in use
Queue counters:      Queued packets  Transmitted packets     Dropped packets
    Low priority                0         336342805223          7986622358
    High priority                0                    0                   0
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets     Dropped packets
  0 best-effort                 0         333760130103                   0
  1 assured-forw                0                    0                   0
  2 mcast-be                    0            274948977                   0
  3 queue3                      0                    0                   0
  4 mcast-ef                    0                    0                   0
  5 expedited-fo                0                    0                   0
  6 mcast-af                    0                    0                   0
  7 network-cont                0                46613                   0
PFE Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets     Dropped packets
  0 best-effort                 0         737867061290          5595302082
  1 assured-forw                0                    0                   0
  2 mcast-be                    0                    0                   0
  3 queue3                      0                    0                   0
  4 mcast-ef                    0                    0                   0
  5 expedited-fo                0                    0                   0
  6 mcast-af                    0                    0                   0
  7 network-cont                0                97800                   0
Active alarms  : None
Active defects : None
MAC statistics:                     Receive         Transmit
  Total octets               49625962140160   47686985710805
  Total packets                387702829264     372554570944
  Unicast packets              387702829264     372554518472
  Broadcast packets                       0                2
  Multicast packets                       0            52470
  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
  Code violations                         0
Packet Forwarding Engine configuration:
  Destination slot: 6
CoS information:
  Direction : Output
  CoS transmit queue            Bandwidth            Buffer Priority   Limit
                      %               bps     %        usec
  0 best-effort        75      7500000000    75           0       low    none
  2 mcast-be           20      2000000000    20           0       low    none
  7 network-cont        5       500000000     5           0       low    none

Logical interface xe-6/0/39.0 (Index 1810) (SNMP ifIndex 2238) (Generation 1923)
```

```
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
 Input  bytes  :                 0
 Output bytes  :           9375416
 Input  packets:                 0
 Output packets:             48901
Local statistics:
 Input  bytes  :                 0
 Output bytes  :           9375416
 Input  packets:                 0
 Output packets:             48901
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: 1937, Route table: 0
  Flags: Trunk-Mode
```

show interfaces queue

Syntax

show interfaces queue
<both-ingress-egress>
<egress>
<forwarding-class *forwarding-class*>
<ingress>
<*interface-name*>

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display class-of-service (CoS) queue information for physical interfaces.

Options

none—Show detailed CoS queue statistics for all physical interfaces.

both-ingress-egress—(Optional) Show both ingress and egress queue statistics. (Ingress statistics are not available for all interfaces.)

egress—(Optional) Show egress queue statistics only.

forwarding-class ***forwarding-class***—(Optional) Show queue statistics only for the specified forwarding class.

ingress—(Optional) Show ingress queue statistics only. (Ingress statistics are not available for all interfaces.)

interface-name—(Optional) Show queue statistics for the specified interface.

Required Privilege Level view

Related Documentation

- Monitoring Interface Status and Traffic on page 289
- *Monitoring Interfaces That Have CoS Components*
- *Defining CoS Schedulers and Scheduler Maps (CLI Procedure)*
- *Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)*

List of Sample Output

show interfaces queue ge-0/0/0 (EX2200 Switch) on page 391
show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch) on page 392

Output Fields Table 50 on page 389 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 50: show interfaces queue Output Fields

Field Name	Field Description
Physical Interface and Forwarding Class Information	
Physical interface	Name of the physical interface.

Table 50: show interfaces queue Output Fields *(continued)*

Field Name	Field Description
Enabled	State of the interface. Possible values are: • **Administratively down, Physical link is Down**—The interface is turned off, and the physical link is inoperable. • **Administratively down, Physical link is Up**—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. • **Enabled, Physical link is Down**—The interface is turned on, but the physical link is inoperable and cannot pass packets. • **Enabled, Physical link is Up**—The interface is turned on, and the physical link is operational and can pass packets.
Interface index	Index number of the physical interface, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Description	User-configured interface description.
Forwarding classes	Number of forwarding classes supported and in use for the interface.
Ingress Queues Information (not shown for all interfaces)	
Ingress queues	Number of input queues supported and in use on the specified interface. For an interface on a line card with oversubscribed ports, the ingress queue handles low priority traffic on the interface.
Transmitted	Transmission statistics for the queue: • **Packets**—Number of packets transmitted by this queue. • **Bytes**—Number of bytes transmitted by this queue. • **Tail-dropped packets**—Number of packets dropped because the queue buffers were full.
PFE chassis queues	For an interface on a line card with oversubscribed ports, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The Packet Forwarding Engine chassis queue for a port group handles high priority traffic from all the interfaces in the port group.
Egress Queues Information	
Egress queues	Number of output queues supported and in use on the specified interface.
Queue	CoS queue number.
Queued	This counter is not supported on EX Series switches.

Table 50: show interfaces queue Output Fields *(continued)*

Field Name	Field Description

Transmitted	Number of packets and bytes transmitted by this queue. Information on transmitted packets and bytes can include:
	• **Packets**—Number of packets transmitted.
	• **Bytes**—Number of bytes transmitted.
	• **Tail-dropped packets**—Number of arriving packets dropped because output queue buffers were full.
	• **RED-dropped packets**—Number of packets dropped because of random early detection (RED).
	  ◦ **Low**—Number of low loss priority packets dropped because of RED.
	  ◦ **High**—Number of high loss priority packets dropped because of RED.
	• **RED-dropped bytes**—Number of bytes dropped because of random early detection (RED).
	  ◦ **Low**—Number of low loss priority bytes dropped because of RED.
	  ◦ **High**—Number of high loss priority bytes dropped because of RED.

Packet Forwarding Engine Chassis Queues	For an interface on a line card with oversubscribed ports, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The queue statistics reflect the traffic flowing on all the interfaces in the port group.

Sample Output

show interfaces queue ge-0/0/0 (EX2200 Switch)

```
user@switch> show interfaces queue ge–0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 130, SNMP ifIndex: 501
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets              :                    0
    Bytes                :                    0
    Tail-dropped packets :                    0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets              :                    0
    Bytes                :                    0
    Tail-dropped packets :                    0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:
    Packets              :                    0
    Bytes                :                    0
    Tail-dropped packets :                    0
Queue: 7, Forwarding classes: network-control
  Queued:
  Transmitted:
    Packets              :                    0
```

```
    Bytes               :                    0
    Tail-dropped packets :                   0
```

show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch)

user@switch> **show interfaces queue xe-6/0/39**

```
Physical interface: xe-6/0/39, Enabled, Physical link is Up
  Interface index: 291, SNMP ifIndex: 1641
Forwarding classes: 16 supported, 7 in use
Ingress queues: 1 supported, 1 in use
  Transmitted:
    Packets             :         337069086018
    Bytes               :       43144843010304
    Tail-dropped packets :          8003867575
PFE chassis queues: 1 supported, 1 in use
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
Forwarding classes: 16 supported, 7 in use
Egress queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets             :         334481399932
    Bytes               :       44151544791024
    Tail-dropped packets :                   0
  Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
  Queue: 2, Forwarding classes: mcast-be
  Queued:
  Transmitted:
    Packets             :            274948977
    Bytes               :          36293264964
    Tail-dropped packets :                   0
Queue: 4, Forwarding classes: mcast-ef
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
 Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
  Queue: 6, Forwarding classes: mcast-af
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
  Queue: 7, Forwarding classes: network-control
  Queued:
  Transmitted:
```



```
    Packets             :                46714
    Bytes               :              6901326
    Tail-dropped packets :                   0

Packet Forwarding Engine Chassis Queues:
Queues: 8 supported, 7 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets             :         739338141426
    Bytes               :       94635282101928
    Tail-dropped packets :                   0
    RED-dropped packets  :          5606426444
     Low                :          5606426444
     High               :                   0
    RED-dropped bytes    :        683262846464
     Low                :        683262846464
     High               :                   0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
    RED-dropped packets  :                   0
     Low                :                    0
     High               :                    0
    RED-dropped bytes    :                   0
     Low                :                    0
     High               :                    0
Queue: 2, Forwarding classes: mcast-be
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
    RED-dropped packets  :                   0
     Low                :                    0
     High               :                    0
    RED-dropped bytes    :                   0
     Low                :                    0
     High               :                    0
Queue: 4, Forwarding classes: mcast-ef
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
    RED-dropped packets  :                   0
     Low                :                    0
     High               :                    0
    RED-dropped bytes    :                   0
     Low                :                    0
     High               :                    0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
    RED-dropped packets  :                   0
```



```
      Low               :                    0
      High              :                    0
     RED-dropped bytes    :                    0
      Low               :                    0
      High              :                    0
Queue: 6, Forwarding classes: mcast-af
  Queued:
  Transmitted:
    Packets             :                    0
    Bytes               :                    0
    Tail-dropped packets :                   0
    RED-dropped packets  :                   0
     Low                :                    0
     High               :                    0
    RED-dropped bytes    :                   0
     Low                :                    0
     High               :                    0
Queue: 7, Forwarding classes: network-control
  Queued:
  Transmitted:
    Packets             :                97990
    Bytes               :             14987506
    Tail-dropped packets :                   0
    RED-dropped packets  :                   0
     Low                :                    0
     High               :                    0
    RED-dropped bytes    :                   0
     Low                :                    0
     High               :                    0
```

show interfaces vlan

Syntax

show interfaces (vlan | vlan.*vlan-id*)
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance (all | *instance-name*)>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display status information about routed VLAN interfaces (RVIs).

Options

vlan | vlan.*vlan-id*—Display status information for the specified RVI.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

routing-instance (all | *instance-name*)—(Optional) Associate this RVI with the named routing instance.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- *show ethernet-switching table*
- *show vlans*
- Monitoring Interface Status and Traffic on page 289
- Troubleshooting Network Interfaces on EX3200 Switches on page 430
- Troubleshooting Network Interfaces on EX4200 Switches on page 432
- *Verifying Routed VLAN Interface Status and Statistics*

List of Sample Output

show interfaces vlan on page 403
show interfaces vlan terse on page 403
show interfaces vlan extensive on page 404
show interfaces vlan detail on page 405

Output Fields

Table 51 on page 396 lists the output fields for the **show interfaces vlan** command. Output fields are listed in the approximate order in which they appear. The level of output *none* means the basic command with no optional options—that is, either just **show interfaces vlan** or **show interfaces vlan.*vlan-id***.

Table 51: show interfaces vlan Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface, which is always **vlan**.	All levels
Enabled	State of the interface: **Enabled** or **Disabled**, followed by the statement **Physical link is <Up/Down>**	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail** **extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail** **extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail** **extensive**
Type	Because this is routed VLAN interface information, this entry is always **VLAN**.	**detail** **extensive** none
Link-level type	Encapsulation (added control information) being used on the physical interface. Because this is routed VLAN interface information, this entry is always **VLAN**.	All levels
MTU	Maximum transmission unit (MTU) size on the physical interface. The default MTU size depends on the switch platform. Changing either the media MTU or protocol MTU causes an interface to be deleted and added again.	All levels
Clocking	Value is always **Unspecified**—not applicable on switches.	**detail** **extensive**
Speed	Speed of the interface, either **Auto** if autonegotiation of speed is enabled or a number representing the configured speed in megabits per second.	**detail** **extensive** none

Table 51: show interfaces vlan Output Fields *(continued)*

Field Name	Field Description	Level of Output
Device flags	Information about the physical device such as: **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status. **Down**—Device has been administratively disabled. **Hear-Own-Xmit**—Device receives its own transmissions. **Is-Default**—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities. **Is-Preferred**—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet. **Is-Primary**—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface. **Link-Layer-Down**—The link-layer protocol has failed to connect with the remote endpoint. **Loopback**—Switch is in physical loopback. **Loop-Detected**—The link layer has received frames that it sent, thereby detecting a physical loopback. **No-Carrier**—On media that support carrier recognition, no carrier is currently detected. **No-Multicast**—Device does not support multicast traffic. **Preferred**—This address is a candidate to become the preferred address. **Present**—Device is physically present and recognized. **Promiscuous**—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. **Primary**—This address is a candidate to become the primary address. **Quench**—Transmission on the device is quenched, because the output buffer is overflowing **Recv-All-Multicasts**—Device is in multicast promiscuous mode and therefore provides no multicast filtering. **Running**—Device is active and enabled.	**detail** **extensive** none
Link type	Link mode of the interface—**Auto** if autonegotiation is enabled, or the configured **Full-Duplex** or **Half-Duplex**.	**detail** **extensive** none
Link flags	Value is always **None**—not applicable on switches.	**detail** **extensive** none
Physical Info	Value is always **Unspecified**—not applicable on switches.	**detail** **extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail** **extensive**
Current address	MAC address of the hardware.	**detail** **extensive** none
Hardware address	MAC address of the switch.	**detail** **extensive** none

Table 51: show interfaces vlan Output Fields *(continued)*

Field Name	Field Description	Level of Output
Alternate link address	Value is always **Unspecified**—not applicable on switches.	**detail** **extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).** The entry can also be **Never**.	**detail** **extensive** none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail** **extensive** none
Traffic statistics	Number and rate of bytes and packets transmitted or received on the physical interface for supported switches. • **Input bytes**—Number of bytes received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter on EX3200 switches and EX4200 switches. EX8200 switches can also be configured to collect this information with the command *l3-interface-ingress-counting*. • **Output bytes**—Number of bytes sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • **Input packets**—Number of packets received on the interface for this switch. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command *l3-interface-ingress-counting*. • **Output packets**—Number of packets sent on the interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches.	**detail** **extensive**
IPv6 transit statistics	Number and rate of bytes and packets transmitted and/or received on the IPv6 interface for supported switches. • **Input bytes**—Number of bytes received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command *l3-interface-ingress-counting*. • **Output bytes**—Number of bytes sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for EX8200 switches. • **Input packets**—Number of packets received on the interface. This value reflects the information gathered by the automatic ingress counter for EX3200 and EX4200 switches. EX8200 switches can also be configured to collect this information with the command *l3-interface-ingress-counting*. • **Output packets**—Number of packets sent on the IPv6 interface. This value reflects the information gathered by the automatic egress counter for and EX8200 switches.	**detail** **extensive**

Table 51: show interfaces vlan Output Fields *(continued)*

Field Name	Field Description	Level of Output
Input Errors	Input errors on the interface. The following paragraphs explain some of the counters whose meaning may not be obvious. • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this value increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—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. • **L3 incompletes**—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. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC. If this value is ever nonzero, the interface is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 51: show interfaces vlan Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This value 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 interface is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—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. • **Collisions**—Number of Ethernet collisions. Both Gigabit Ethernet interfaces and 10 Gigabit Ethernet interfaces support only full-duplex operation, so for those two interfaces, this value should always be zero. If the value is nonzero for either Gigabit Ethernet or 10 Gigabit Ethernet, there is a software bug. • **Aged packets**—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. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the interface. If this value is ever nonzero, the interface is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
vlan.*vlan-id*, Index, SNMP ifIndex	VLAN ID, index, and SNMP index number for the logical interface. The logical interface index values reflect the item's initialization sequence.	**detail** **extensive** none
Generation	Unique number for Juniper Networks Technical support use only.	**detail** **extensive** none
Flags	Errors that have occurred on this interface, such as **Link Layer Down**. Other possible flags include: • **Device-down**—Device has been administratively disabled. • **Disabled**—Interface is administratively disabled. • **Down**—A hardware failure has occurred. • **Hardware-Down**—Interface protocol initialization failed to complete successfully. • **SNMP-Traps**—SNMP trap notifications are enabled. • **Up**—Interface is enabled and operational.	**detail** **extensive** none

Table 51: show interfaces vlan Output Fields *(continued)*

Field Name	Field Description	Level of Output
SNMP-Traps	Each configured SNMP trap has a number that appears here—0x0 is always displayed for logical interface SNMP traps.	**detail** **extensive** none
Encapsulation	Encapsulation method, which is the process of adding control information. The value is always Ethernet 2 (**ENET2**) for logical encapsulation.	**detail** **extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical interface of supported switches. Traffic statistics represent the sum of the next two fields, Local statistics and Transit statistics. Note that these are not the values for the RVI ingress or egress counters—for that value, see Transit statistics below. • **Input bytes**—Number of bytes received on the interface. Same value as the physical interface. • **Output bytes**—Number of bytes sent on the interface. Same value as the physical interface. • **Input packets**—Number of packets received on the interface. Same value as the physical interface. • **Output packets**—Number of bytes sent on the interface. Same value as the physical interface. NOTE: The bandwidth bps counter is not enabled on the switches.	**detail** **extensive**
Local statistics	Number and rate of bytes and packets received and transmitted locally by the Routing Engine on the logical interface of supported switches. All packets for protocols and process statistics are counted here. • **Input bytes**—Number of bytes received on the interface. Same value as for the physical interface. • **Output bytes**—Number of bytes sent on the interface. Same value as for the physical interface. • **Input packets**—Number of packets received on the interface. Same value as for the physical interface. • **Output packets**—Number of bytes sent on the interface. Same value as for the physical interface.	**detail** **extensive** none
Transit statistics	Number and rate of bytes and packets received and transmitted on the RVI logical interface of supported switches. Look at this value to see the RVI ingress and egress count. • **Input bytes**—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • **Output bytes**—Number of bytes sent on the interface. This egress counter is automatic for EX8200. • **Input packets**—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • **Output packets**—Number of packets sent on the interface. This egress counter is automatic for EX8200 switches.	**detail** **extensive**

Table 51: show interfaces vlan Output Fields *(continued)*

Field Name	Field Description	Level of Output
IPv6 transit statistics	Number and rate of IPv6 bytes and packets received and transmitted on the RVI logical interface of supported switches. Transit values are unique to the logical interface and do not appear in physical interface output. Look at the values listed below to see the RVI ingress and egress count for IPv6 traffic. • **Input bytes**—Number of bytes received on the interface. This ingress counter is automatic for EX3200 and EX4200 switches and configurable for EX8200 switches. • **Output bytes**—Number of bytes sent by the interface. This egress counter is automatic for EX8200 switches. • **Input packets**—Number of packets received on the interface. This ingress counter is automatic for EX3200 and EX4200 and configurable for EX8200 switches. • **Output packets**—Number of packets sent by the interface. This egress counter is automatic for EX8200 switches. NOTE: The bandwidth bps counter is not enabled on the switches.	**detail** **extensive**
Protocol	Protocol used for the logical interface—this value is **inet** for IPv4 traffic and **inet6** for IPv6 traffic.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	**detail** **extensive**
Route table	Route table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail** **extensive** none
Protocol flags	Information about the protocol such as **Targeted-broadcast**.	**detail** **extensive** none
Protocol addresses and Address flags	Protocol address values here can be: **Dest-route-down**—The routing process detected that the link was not operational and changed the interface routes to nonforwarding status **Device-down**—Device has been administratively disabled. **Disabled**—Interface is administratively disabled. **Down**—A hardware failure has occurred. **Hardware-Down**—Interface protocol initialization failed to complete successfully. **Is-Default**—This address is the default address of the switch. The default address is used as the source address by SNMP, ping, traceroute, and other network utilities. **Is-Preferred**—This address is the default local address for packets originating from the local switch and sent to destinations on the subnet. **Is-Primary**—This address is the default local address for broadcast and multicast packets originated locally and sent out the interface. **Preferred**—This address is a candidate to become the preferred address. **Primary**—This address is a candidate to become the primary address. **SNMP-Traps**—SNMP trap notifications are enabled. **Up**—Interface is enabled and operational.	**detail** **extensive** none

Table 51: show interfaces vlan Output Fields *(continued)*

Field Name	Field Description	Level of Output
Address destination	Logical destination's network address.	**detail** **extensive** none
Local address	IP address of the logical interface.	**detail** **extensive** none
Broadcast address	Broadcast address of the logical interface.	**detail** **extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail** **extensive**

Sample Output

show interfaces vlan

```
user@switch> show interfaces vlan
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface vlan.0 (Index 82) (SNMP ifIndex 557)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 1
    Protocol inet
      Flags: Targeted-broadcast
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255

  Logical interface vlan.1 (Index 83) (SNMP ifIndex 558)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Input packets : 0
    Output packets: 1
    Protocol inet
      Flags: Targeted-broadcast
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255
```

show interfaces vlan terse

```
user@switch> show interfaces vlan terse
Interface               Admin Link Proto    Local                 Remote
vlan                    up    up
```

```
vlan.0                  up    down inet     10.1.1.1/24
vlan.1                  up    down inet     10.1.2.1/24
```

show interfaces vlan extensive

```
user@switch> show interfaces vlan extensive
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556, Generation: 153
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
   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, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                    0
     Output bytes  :                   42
     Input  packets:                    0
     Output packets:                    1
    Local statistics:
     Input  bytes  :                    0
     Output bytes  :                   42
     Input  packets:                    0
     Output packets:                    1
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, Generation: 159, Route table: 0
      Flags: Targeted-broadcast
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
        Generation: 138

  Logical interface vlan.1 (Index 83) (SNMP ifIndex 558) (Generation 148)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
```

```
      Traffic statistics:
       Input  bytes  :                   0
       Output bytes  :                  42
       Input  packets:                   0
       Output packets:                   1
      Local statistics:
       Input  bytes  :                   0
       Output bytes  :                  42
       Input  packets:                   0
       Output packets:                   1
      Transit statistics:
       Input  bytes  :                   0                    0 bps
       Output bytes  :                   0                    0 bps
       Input  packets:                   0                    0 pps
      Output packets:                   0                    0 pps
     Protocol inet, Generation: 160, Route table: 0
       Flags: Targeted-broadcast
       Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
         Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255,
         Generation: 140
```

show interfaces vlan detail

```
user@switch> show interfaces vlan detail
Physical interface: vlan, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 556, Generation: 153
  Type: VLAN, Link-level type: VLAN, MTU: 1518, Clocking: Unspecified,
  Speed: 1000mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:21:59:c5:f0:40, Hardware address: 00:21:59:c5:f0:40
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0

  Logical interface vlan.0 (Index 82) (SNMP ifIndex 557) (Generation 147)
    Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                  42
     Input  packets:                   0
     Output packets:                   1
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                  42
     Input  packets:                   0
     Output packets:                   1
```

```
  Transit statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
  Protocol inet, Generation: 159, Route table: 0
    Flags: Targeted-broadcast
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
      Generation: 138

Logical interface vlan.1 (Index 83) (SNMP ifIndex 558) (Generation 148)
  Flags: Link-Layer-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                  42
   Input  packets:                   0
   Output packets:                   1
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                  42
   Input  packets:                   0
   Output packets:                   1
  Transit statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
  Protocol inet, Generation: 160, Route table: 0
    Flags: Targeted-broadcast
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.1.2/24, Local: 10.1.2.1, Broadcast: 10.1.2.255,
      Generation: 140
```

show lacp interfaces

Syntax

show lacp interfaces
<*interface-name*>

Release Information

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.

Description

Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.

Options

none—Display LACP information for all interfaces.

interface-name—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*x***
- Gigabit Ethernet—**ge-*fpc/pic/port***
- 10-Gigabit Ethernet—**xe-*fpc/pic/port***

Required Privilege Level

view

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 121
- *Configuring Link Aggregation*
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 125
- *Configuring Aggregated Ethernet LACP*
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 126
- Understanding Aggregated Ethernet Interfaces and LACP on page 8
- *Understanding Aggregated Ethernet Interfaces and LACP*
- *Junos OS Interfaces Fundamentals Configuration Guide*

List of Sample Output

show lacp interfaces (EX Series Switches) on page 409
show lacp interfaces (QFX Series) on page 410

Output Fields Table 52 on page 408 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 52: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated Ethernet interface name.

LACP State	LACP state information for each aggregated Ethernet interface:
	• For a child interface configured with the **force-up** statement, LACP state displays **FUP** along with the interface name.
	• **Role**—Role played by the interface. It can be one of the following:
	  • **Actor**—Local device participating in the LACP negotiation.
	  • **Partner**—Remote device participating in the LACP negotiation.
	• **Exp**—Expired state. **Yes** indicates that the actor or partner is in an expired state. **No** indicates that the actor or partner is not in an expired state.
	• **Def**—Default. **Yes** indicates that the actor's receive machine is using the default operational partner information, which is administratively configured for the partner. **No** indicates that the operational partner information in use has been received in an LACP PDU.
	• **Dist**—Distribution of outgoing frames. **No** indicates that the distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is **Yes**.
	• **Col**—Collection of incoming frames. **Yes** indicates that the collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is **No**.
	• **Syn**—Synchronization. If the value is **Yes**, 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 **No**, the link is not synchronized. The link is currently not in the right aggregation.
	• **Aggr**—Ability of the aggregation port to aggregate (**Yes**) or to operate only as an individual link (**No**).
	• **Timeout**—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 (**Long Timeout** or **Short Timeout**).
	• **Activity**—Actor's or partner's port activity. **Passive** indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is **Active**. **Active** indicates the port's preference to participate in the protocol regardless of the partner's control value.

Table 52: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description
LACP Protocol	LACP protocol information for each aggregated interface: • Link state (active or standby) indicated in parentheses next to the interface when link protection is configured. • **Receive State**—One of the following values: ◦ **Current**—The state machine receives an LACP PDU and enters the **Current** state. ◦ **Defaulted**—If no LACP PDU is received before the timer for the **Current** state expires a second time, the state machine enters the **Defaulted** state. ◦ **Expired**—If no LACP PDU is received before the timer for the **Current** state expires once, the state machine enters the **Expired** state. ◦ **Initialize**—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the **Initialize** state. ◦ **LACP Disabled**—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to **LACP Disabled**. This state is similar to the **Defaulted** state, except that the port is forced to operate as an individual port. ◦ **Port Disabled**—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the **Port Disabled** state. • **Transmit State**—Transmit state of the state machine. The transmit state is one of the following values: ◦ **Fast periodic**—Periodic transmissions are enabled at a fast transmission rate. ◦ **No periodic**—Periodic transmissions are disabled. ◦ **Periodic timer**—Transitory state entered when the periodic timer expires. ◦ **Slow periodic**—Periodic transmissions are enabled at a slow transmission rate. • **Mux State**—State of the multiplexer state machine for the aggregation port. The state is one of the following values: ◦ **Attached**—The multiplexer state machine initiates the process of attaching the port to the selected aggregator. ◦ **Collecting**—**Yes** 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. **No** indicates the receive function of this link is not enabled. ◦ **Collecting distributing**—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. ◦ **Detached**—Process of detaching the port from the aggregator is in progress. ◦ **Distributing**—**Yes** 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. **No** indicates the transmit function of this link is not enabled. ◦ **Waiting**—The multiplexer state machine is in a holding process, awaiting an outcome.

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 virtual-chassis vc-port diagnostics optics

Syntax

show virtual-chassis vc-port diagnostics optics
<all-members>
<*interface-name*>
<local>
<member *member-id*>

Release Information

Command introduced in Junos OS Release 12.2 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D20 for Virtual Chassis Fabric (VCF).

Description

Display diagnostics data and alarms for Ethernet optical transceivers installed in ports configured as Virtual Chassis Ports (VCPs) in an EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.

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 a transceiver is not operating properly. DOM information can be used to diagnose why a transceiver is not working.

On some EX Series switches, the **request virtual-chassis vc-port diagnostics optics** command must be entered to run a diagnostic scan before you can gather the **show virtual-chassis vc-port diagnostics optics** output.

Options

none—Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.

all-members—(Optional) Display diagnostics information for transceivers installed in VCPs of all members of a Virtual Chassis or VCF.

interface-name—(Optional) Display diagnostics information for the transceiver installed in a specified VCP.

local—(Optional) Display diagnostics information for transceivers installed in VCPs on the switch or external Routing Engine on which this command is entered.

member *member-id*—(Optional) Display diagnostics information for transceivers installed in VCPs on a specified member of a Virtual Chassis or VCF.

Required Privilege Level

view

Related Documentation

- *show virtual-chassis vc-port*
- *Installing a Transceiver in an EX Series Switch*
- *Removing a Transceiver from a Switch*
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

show virtual-chassis vc-port diagnostics optics on page 415
show virtual-chassis vc-port diagnostics optics (interface-name) on page 420

show virtual-chassis vc-port diagnostics optics local on page 422
show virtual-chassis vc-port diagnostics optics (member member-id) on page 424

Output Fields Table 53 on page 413 lists the output fields for the **show virtual-chassis vc-port diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

Table 53: show virtual-chassis vc-port diagnostics optics Output Fields

Field Name	Field Description
FPC	Displays the FPC slot number.
Virtual chassis port	Displays the name of the VCP.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes (mA). The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is *On* or *Off*.
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is *On* or *Off*.
Laser bias current high warning	Displays whether the laser bias power setting high warning is *On* or *Off*.
Laser bias current low warning	Displays whether the laser bias power setting low warning is *On* or *Off*.
Laser output power high alarm	Displays whether the laser output power high alarm is *On* or *Off*.
Laser output power low alarm	Displays whether the laser output power low alarm is *On* or *Off*.
Laser output power high warning	Displays whether the laser output power high warning is *On* or *Off*.
Laser output power low warning	Displays whether the laser output power low warning is *On* or *Off*.
Module temperature high alarm	Displays whether the module temperature high alarm is *On* or *Off*.
Module temperature low alarm	Displays whether the module temperature low alarm is *On* or *Off*.
Module temperature high warning	Displays whether the module temperature high warning is *On* or *Off*.
Module temperature low warning	Displays whether the module temperature low warning is *On* or *Off*.

Table 53: show virtual-chassis vc-port diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Module voltage high alarm	Displays whether the module voltage high alarm is *On* or *Off*.
Module voltage low alarm	Displays whether the module voltage low alarm is *On* or *Off*.
Module voltage high warning	Displays whether the module voltage high warning is *On* or *Off*.
Module voltage low warning	Displays whether the module voltage low warning is *On* or *Off*.
Laser rx power high alarm	Displays whether the receive laser power high alarm is *On* or *Off*.
Laser rx power low alarm	Displays whether the receive laser power low alarm is *On* or *Off*.
Laser rx power high warning	Displays whether the receive laser power high warning is *On* or *Off*.
Laser rx power low warning	Displays whether the receive laser power low warning is *On* or *Off*.
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.

Table 53: show virtual-chassis vc-port diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

Sample Output

show virtual-chassis vc-port diagnostics optics

```
user@switch> show virtual-chassis vc-port diagnostics optics
fpc0:
--------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                        :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                        :  N/A

fpc1:
--------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                        :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                        :  N/A

fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-2/0
    Optical diagnostics                        :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                        :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                        :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                        :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                         :  4.130 mA
    Laser output power                         :  0.2450 mW / -6.11 dBm
    Module temperature                         :  32 degrees C / 90 degrees F
    Module voltage                             :  3.3530 V
    Receiver signal average optical power      :  0.0971 mW / -10.13 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   :  14.998 mA
        Laser bias current low alarm threshold    :  0.998 mA
        Laser bias current high warning threshold :  14.000 mA
        Laser bias current low warning threshold  :  1.198 mA
        Laser output power high alarm threshold   :  0.7940 mW / -1.00 dBm
        Laser output power low alarm threshold    :  0.0790 mW / -11.02 dBm
        Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
        Laser output power low warning threshold  :  0.0990 mW / -10.04 dBm
        Module temperature high alarm threshold   :  85 degrees C / 185 degrees F
        Module temperature low alarm threshold    :  -10 degrees C / 14 degrees F
        Module temperature high warning threshold :  80 degrees C / 176 degrees F
        Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
        Module voltage high alarm threshold      :  3.600 V
        Module voltage low alarm threshold       :  3.000 V
        Module voltage high warning threshold    :  3.499 V
        Module voltage low warning threshold     :  3.099 V
        Laser rx power high alarm threshold      :  1.5848 mW / 2.00 dBm
        Laser rx power low alarm threshold       :  0.0100 mW / -20.00 dBm
        Laser rx power high warning threshold    :  1.2589 mW / 1.00 dBm
        Laser rx power low warning threshold     :  0.0125 mW / -19.03 dBm
    Virtual chassis port: vcp-255/0/3
        Laser bias current                       :  5.428 mA
        Laser output power                       :  0.4760 mW / -3.22 dBm
        Module temperature                       :  28 degrees C / 83 degrees F
        Module voltage                           :  3.3440 V
        Receiver signal average optical power    :  0.4002 mW / -3.98 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   :  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

fpc3:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/2
    Laser bias current                        :  7.876 mA
    Laser output power                        :  0.5330 mW / -2.73 dBm
    Module temperature                        :  26 degrees C / 78 degrees F
    Module voltage                            :  3.3060 V
    Receiver signal average optical power     :  0.4885 mW / -3.11 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   :  14.500 mA
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 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   :  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       :  1.9952 mW / 3.00 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
    Virtual chassis port: vcp-255/0/3
        Laser bias current                        :  5.052 mA
        Laser output power                        :  0.5030 mW / -2.98 dBm
        Module temperature                        :  24 degrees C / 75 degrees F
        Module voltage                            :  3.2890 V
        Receiver signal average optical power     :  0.5028 mW / -2.99 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   :  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
    Virtual chassis port: vcp-255/0/4
        Laser bias current                        :  7.978 mA
        Laser output power                        :  0.5460 mW / -2.63 dBm
        Module temperature                        :  24 degrees C / 76 degrees F
```

```
    Module voltage                            :  3.3060 V
    Receiver signal average optical power     :  0.6305 mW / -2.00 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   :  14.500 mA
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 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   :  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       :  1.9952 mW / 3.00 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
fpc4:
--------------------------------------------------------------------------
Virtual chassis port: vcp-0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/4
    Laser bias current                        :  7.860 mA
    Laser output power                        :  0.5370 mW / -2.70 dBm
    Module temperature                        :  24 degrees C / 75 degrees F
    Module voltage                            :  3.2920 V
    Receiver signal average optical power     :  0.6271 mW / -2.03 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
    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   :  14.500 mA
    Laser bias current low alarm threshold    :  3.500 mA
    Laser bias current high warning threshold :  14.500 mA
    Laser bias current low warning threshold  :  3.500 mA
    Laser output power high alarm threshold   :  1.8620 mW / 2.70 dBm
    Laser output power low alarm threshold    :  0.0740 mW / -11.31 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   :  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       :  1.9952 mW / 3.00 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
```

show virtual-chassis vc-port diagnostics optics (interface-name)

```
user@external-routing-engine> show virtual-chassis vc-port diagnostics optics vcp-255/0/3
fpc0:
--------------------------------------------------------------------------

fpc1:
--------------------------------------------------------------------------

fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.448 mA
    Laser output power                        :  0.4770 mW / -3.21 dBm
    Module temperature                        :  28 degrees C / 82 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.3973 mW / -4.01 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   :  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

fpc3:
--------------------------------------------------------------------------
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.040 mA
    Laser output power                        :  0.5020 mW / -2.99 dBm
    Module temperature                        :  24 degrees C / 74 degrees F
    Module voltage                            :  3.2870 V
    Receiver signal average optical power     :  0.5073 mW / -2.95 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   :  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

fpc4:
--------------------------------------------------------------------------
```

show virtual-chassis vc-port diagnostics optics local

```
user@switch> show virtual-chassis vc-port diagnostics optics local
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                        :  4.130 mA
    Laser output power                        :  0.2450 mW / -6.11 dBm
    Module temperature                        :  32 degrees C / 90 degrees F
    Module voltage                            :  3.3530 V
    Receiver signal average optical power     :  0.0961 mW / -10.17 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   :  14.998 mA
    Laser bias current low alarm threshold    :  0.998 mA
```

```
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  1.198 mA
    Laser output power high alarm threshold   :  0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold    :  0.0790 mW / -11.02 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold   :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold    :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  80 degrees C / 176 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold       :  3.600 V
    Module voltage low alarm threshold        :  3.000 V
    Module voltage high warning threshold     :  3.499 V
    Module voltage low warning threshold      :  3.099 V
    Laser rx power high alarm threshold       :  1.5848 mW / 2.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold      :  0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.426 mA
    Laser output power                        :  0.4760 mW / -3.22 dBm
    Module temperature                        :  28 degrees C / 83 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.3955 mW / -4.03 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
    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   :  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
```

show virtual-chassis vc-port diagnostics optics (member member-id)

```
user@switch> show virtual-chassis vc-port diagnostics optics member 2
fpc2:
--------------------------------------------------------------------------
Virtual chassis port: vcp-2/0
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-2/1
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/14
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/15
    Optical diagnostics                       :  N/A
Virtual chassis port: vcp-255/0/24
    Laser bias current                        :  4.130 mA
    Laser output power                        :  0.2450 mW / -6.11 dBm
    Module temperature                        :  31 degrees C / 88 degrees F
    Module voltage                            :  3.3530 V
    Receiver signal average optical power     :  0.0961 mW / -10.17 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   :  14.998 mA
    Laser bias current low alarm threshold    :  0.998 mA
    Laser bias current high warning threshold :  14.000 mA
    Laser bias current low warning threshold  :  1.198 mA
    Laser output power high alarm threshold   :  0.7940 mW / -1.00 dBm
    Laser output power low alarm threshold    :  0.0790 mW / -11.02 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0990 mW / -10.04 dBm
    Module temperature high alarm threshold   :  85 degrees C / 185 degrees F
    Module temperature low alarm threshold    :  -10 degrees C / 14 degrees F
    Module temperature high warning threshold :  80 degrees C / 176 degrees F
    Module temperature low warning threshold  :  -5 degrees C / 23 degrees F
    Module voltage high alarm threshold       :  3.600 V
    Module voltage low alarm threshold        :  3.000 V
    Module voltage high warning threshold     :  3.499 V
    Module voltage low warning threshold      :  3.099 V
    Laser rx power high alarm threshold       :  1.5848 mW / 2.00 dBm
```

```
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  1.2589 mW / 1.00 dBm
    Laser rx power low warning threshold      :  0.0125 mW / -19.03 dBm
Virtual chassis port: vcp-255/0/3
    Laser bias current                        :  5.418 mA
    Laser output power                        :  0.4770 mW / -3.21 dBm
    Module temperature                        :  28 degrees C / 83 degrees F
    Module voltage                            :  3.3450 V
    Receiver signal average optical power     :  0.3964 mW / -4.02 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   :  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
```


test interface restart-auto-negotiation

Syntax test interface restart-auto-negotiation *interface-name*

Release Information Command introduced in Junos OS Release 7.6.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Restarts auto-negotiation on a Fast Ethernet or Gigabit Ethernet interface.

Options *interface-name*—Interface name: **fe-*fpc/pic/port*** or **ge-*fpc/pic/port***.

Required Privilege Level view

List of Sample Output test interface restart-auto-negotiation on page 426

Output Fields Use the **show interfaces extensive** command to see the state for auto-negotiation.

Sample Output

test interface restart-auto-negotiation

`user@host>` **test interface restart-auto-negotiation fe-1/0/0**

PART 4

Troubleshooting

CHAPTER 7

Troubleshooting Procedures

Troubleshooting an Aggregated Ethernet Interface

Troubleshooting issues for aggregated Ethernet interfaces:

Show Interfaces Command Shows the LAG is Down

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).
- 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 (or the same Virtual Chassis).

Logical Interface Statistics Do Not Reflect All Traffic

Problem **Description:** The traffic statistics for a logical interface do not include all of the traffic.

Solution Traffic statistics fields for logical interfaces in **show interfaces** commands show only control traffic; the traffic statistics do not include data traffic. You can view the statistics for all traffic only per physical interface.

IPv6 Interface Traffic Statistics Are Not Supported

Problem **Description:** The **IPv6 transit statistics** in the **show interfaces** command display all **0** values.

Solution EX Series switches do not support the collection and reporting of IPv6 transit statistics.

SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0

Problem **Description:** The values for the SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are always 0.

Solution The SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are not supported for aggregated Ethernet interfaces on EX Series switches.

Related Documentation

- Verifying the Status of a LAG Interface on page 291
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 33
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch on page 39

Troubleshooting Network Interfaces on EX3200 Switches

This topic provides troubleshooting information for specific problems related to interfaces on EX3200 switches.

- The interface on one of the last four built-in network ports in an EX3200 switch (for example, interface ge-0/0/23) is down on page 431
- The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down on page 431

The interface on one of the last four built-in network ports in an EX3200 switch (for example, interface ge-0/0/23) is down

Problem

Description: The interface on one of the last four built-in ports (**ge-0/0/20** through **ge-0/0/23** on 24-port models or **ge-0/0/44** through **ge-0/0/47** on 48-port models) of an EX3200 switch is down.

Environment: An SFP or SFP+ uplink module is installed in the switch and a transceiver is installed in one of the ports on the uplink module.

Symptoms: When you check the status with the CLI command **show interfaces ge-** or with the J-Web user interface, the disabled port is not listed.

Cause

The last four built-in ports use the same ASIC as the SFP uplink module. Therefore, if you install a transceiver in an SFP or SFP+ uplink module installed in an EX3200 switch, a corresponding base port from the last four built-in ports is disabled.

Solution

If you need to use the disabled built-in port, you must remove the transceiver from the SFP or SFP+ uplink module. Alternatively, you can install an XFP uplink module instead of an SFP or SFP+ uplink module. There is no conflict between the built-in network ports and the ports on the XFP uplink modules.

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

Problem

Description: The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module installed in an EX3200 switch is down.

Symptoms: When you check the status with the CLI command **show interfaces ge-** or with the J-Web user interface, the disabled port is not listed.

Cause

By default, the SFP+ uplink module operates in the 10-gigabit mode and supports only SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution

Either SFP+ or SFP transceivers can be installed in SFP+ uplink modules. You must configure the operating mode of the SFP+ uplink module to match the type of transceiver you want to use. For SFP+ transceivers, configure the 10-gigabit operating mode and for SFP transceivers, configure the 1-gigabit operating mode. See “Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)” on page 140.

Related Documentation

- Troubleshooting Uplink Module Installation or Replacement on EX3200 Switches on page 433
- Monitoring Interface Status and Traffic on page 289
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- *Removing a Transceiver from a Switch*
- *Uplink Modules in EX3200 Switches*
- EX Series Switches Interfaces Overview on page 3

Troubleshooting Network Interfaces on EX4200 Switches

This topic provides troubleshooting information for specific problems related to interfaces on EX4200 switches.

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

Problem

Description: The interface on the port in which an SFP or SFP+ transceiver is installed in an uplink module installed in an EX4200 switch is down.

Symptoms: When you check the status with the CLI command **show interfaces ge-** or with the J-Web user interface, the disabled port is not listed.

Cause

By default, the SFP+ and SFP+ MACsec uplink modules operate in the 10-gigabit mode and support only SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution

Either SFP+ or SFP transceivers can be installed in the uplink modules. You must configure the operating mode of the SFP+ or SFP+ MACsec uplink module to match the type of transceiver you want to use. For SFP+ transceivers, configure the 10-gigabit operating mode and for SFP transceivers, configure the 1-gigabit operating mode. See “Setting the Mode on an SFP+ or SFP+ MACSec Uplink Module (CLI Procedure)” on page 140.

Related Documentation

- *Troubleshooting Virtual Chassis Port Connectivity on an EX4200 Switch*
- Monitoring Interface Status and Traffic on page 289
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- *Removing a Transceiver from a Switch*
- *Uplink Modules in EX4200 Switches*
- EX Series Switches Interfaces Overview on page 3

Troubleshooting Uplink Module Installation or Replacement on EX3200 Switches

This topic provides troubleshooting information for specific problems related to uplink module ports on EX3200 switches.

1. One of the last four network ports on an EX3200 switch with an SFP or SFP+ uplink module installed is disabled on page 433
2. A port on an SFP uplink module installed in an EX3200 switch is disabled on page 433

One of the last four network ports on an EX3200 switch with an SFP or SFP+ uplink module installed is disabled

Problem

Description: One of the last four built-in ports (**ge-0/0/20** through **ge-0/0/23** on 24-port models or **ge-0/0/44** through **ge-0/0/47** on 48-port models) of an EX3200 switch with an SFP or SFP+ uplink module installed in it is disabled.

Symptoms: When you check the status with the CLI command **show interfaces ge-** or with the J-Web user interface, the disabled port is not listed.

Cause

The last four built-in ports use the same ASIC as the SFP uplink module. Therefore, if you install a transceiver in an SFP or SFP+ uplink module installed in an EX3200 switch, a corresponding base port from the last four built-in ports is disabled.

Solution

If you need to use the disabled built-in port, you must remove the transceiver from the SFP or SFP+ uplink module. Alternatively, you can install an XFP uplink module instead of an SFP or SFP+ uplink module. There is no conflict between the built-in network ports and the ports on the XFP uplink modules.

A port on an SFP uplink module installed in an EX3200 switch is disabled

Problem

Description: One of the ports (**ge-0/1/0** through **ge-0/1/3**) of an SFP uplink module installed in an EX3200 switch is disabled.

Symptoms: When you check the status with the **show interface** commands or with the J-Web user interface, the disabled port is not listed.

Cause

If you replace a transceiver multiple times in quick succession in a port in an SFP uplink module installed in an EX3200 switch, it might cause an eeprom read problem. The switch might not create an interface for that port and that port might be disabled.

Solution

To enable the disabled uplink module port, remove the transceiver from that port and install it after 10 seconds.

Related Documentation

- Monitoring Interface Status and Traffic on page 289
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- *Installing an Uplink Module in an EX3200 Switch*
- *Removing a Transceiver from a Switch*
- *Uplink Modules in EX3200 Switches*

Troubleshooting Interface Configuration and Cable Faults

NOTE: This topic applies only to the J-Web Application package.

Troubleshooting interface configuration and connectivity on the EX Series switch:

1. Interface Configuration or Connectivity Is Not Working on page 434

Interface Configuration or Connectivity Is Not Working

Problem **Description:**

NOTE: This topic applies only to the J-Web Application package.

You encounter errors when you attempt to configure an interface on the switch, or the interface is exhibiting connectivity problems.

Solution Use the port troubleshooter feature in the J-Web interface to identify and rectify port configuration and connectivity related problems.

To use the J-Web interface port troubleshooter:

1. Select the option **Troubleshoot** from the main menu.
2. Click **Troubleshoot Port**. The Port Troubleshooting wizard is displayed. Click **Next**.
3. Select the ports to troubleshoot.
4. Select the test cases to be executed on the selected port. Click **Next**.

   When the selected test cases are executed, the final result and the recommended action is displayed.

If there is a cable fault, the port troubleshooter displays details and the recommended action. For example, the cable must be replaced.

If the port configuration needs to be modified, the port troubleshooter displays details and the recommended action.

Related Documentation

- Monitoring Interface Status and Traffic on page 289

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 70
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 66
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Connecting and Configuring an EX Series Switch (CLI Procedure)*
- *Connecting and Configuring an EX Series Switch (J-Web Procedure)*

Troubleshooting Unicast RPF

Troubleshooting issues for unicast reverse-path forwarding (RPF) on EX Series switches include:

1. Legitimate Packets Are Discarded on page 435

Legitimate Packets Are Discarded

Problem **Description:** The switch filters valid packets from legitimate sources, which results in the switch's discarding packets that should be forwarded.

Solution The interface or interfaces on which legitimate packets are discarded are asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, so the interface that receives a packet is not the same interface the switch uses to reply to the packet's source.

Unicast RPF works properly only on symmetrically routed interfaces. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Unicast RPF filters packets by checking the forwarding table for the best return path to the source of an incoming packet. If the best return path uses the same interface as the interface that received the packet, the switch forwards the packet. If the best return path uses a different interface than the interface that received the packet, the switch discards the packet.

NOTE: On EX3200, EX4200, and EX4300 switches, unicast RPF works properly only if all switch interfaces—including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs), integrated routing and bridging (IRB) interfaces, and routed VLAN interfaces (RVIs)—are symmetrically routed, because unicast RPF is enabled globally on all switch interfaces.

Related Documentation

- Verifying Unicast RPF Status on page 293
- Understanding Unicast RPF on page 12

Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure)

Problem

Description: A 10/100/1000BASE-T Ethernet interface has connectivity problems that you suspect might be caused by a faulty cable.

Solution

Use the time domain reflectometry (TDR) test to determine whether a twisted-pair Ethernet cable is faulty.

The TDR test:

- Detects and reports faults for each twisted pair in an Ethernet cable. Faults detected include open circuits, short circuits, and impedance mismatches.
- Reports the distance to fault to within 1 meter.
- Detects and reports pair swaps, pair polarity reversals, and excessive pair skew.

The TDR test is supported on the following switches and interfaces:

- EX2200, EX3200, EX3300, and EX4200 switches—RJ-45 network interfaces. The TDR test is not supported on management interfaces and SFP interfaces.
- EX6200 and EX8200 switches—RJ-45 network interfaces on line cards.

NOTE: We recommend running the TDR test on an interface when there is no traffic on the interface.

To diagnose a cable problem by running the TDR test:

1. Run the request diagnostics tdr command.

```
user@switch> request diagnostics tdr start interface ge-0/0/10

Interface TDR detail:
Test status                        : Test successfully executed  ge-0/0/10
```

2. View the results of the TDR test with the show diagnostics tdr command.

```
user@switch> show diagnostics tdr interface ge-0/0/10

Interface TDR detail:
Interface name                     : ge-0/0/10
Test status                        : Passed
Link status                        : Down
MDI pair                           : 1-2
  Cable status                     : Normal
  Distance fault                   : 0 Meters
  Polartiy swap                    : N/A
  Skew time                        : N/A
MDI pair                           : 3-6
  Cable status                     : Normal
  Distance fault                   : 0 Meters
  Polartiy swap                    : N/A
  Skew time                        : N/A
```

```
MDI pair                        : 4-5
  Cable status                  : Open
  Distance fault                : 1 Meters
  Polartiy swap                 : N/A
  Skew time                     : N/A
MDI pair                        : 7-8
  Cable status                  : Normal
  Distance fault                : 0 Meters
  Polartiy swap                 : N/A
  Skew time                     : N/A
Channel pair                    : 1
  Pair swap                     : N/A
Channel pair                    : 2
  Pair swap                     : N/A
Downshift                       : N/A
```

3. Examine the **Cable status** field for the four MDI pairs to determine if the cable has a fault. In the preceding example, the twisted pair on pins 4 and 5 is broken or cut at approximately one meter from the **ge-0/0/10** port connection.

NOTE: The **Test Status** field indicates the status of the TDR test, not the cable. The value **Passed** means the test completed—it does not mean that the cable has no faults.

The following is additional information about the TDR test:

- The TDR test can take some seconds to complete. If the test is still running when you execute the **show diagnostics tdr** command, the **Test status** field displays **Started**. For example:

  ```
  user@switch> show diagnostics tdr interface ge-0/0/22

  Interface TDR detail:
  Interface name                : ge-0/0/22
  Test status                   : Started
  ```

- You can terminate a running TDR test before it completes by using the **request diagnostics tdr abort interface *interface-name*** command. The test terminates with no results, and the results from any previous test are cleared.

- You can display summary information about the last TDR test results for all interfaces on the switch that support the TDR test by not specifying an interface name with the **show diagnostics tdr** command. For example:

  ```
  user@switch> show diagnostics tdr
  Interface    Test status    Link status  Cable status  Max distance fault
   ge-0/0/0         Passed      UP           OK             0
   ge-0/0/1    Not Started     N/A          N/A           N/A
   ge-0/0/2         Passed      UP           OK             0
   ge-0/0/3    Not Started     N/A          N/A           N/A
   ge-0/0/4         Passed      UP           OK             0
   ge-0/0/5         Passed      UP           OK             0
   ge-0/0/6         Passed      UP           OK             0
   ge-0/0/7    Not Started     N/A          N/A           N/A
   ge-0/0/8         Passed     Down          OK             0
   ge-0/0/9    Not Started     N/A          N/A           N/A
  ```

```
ge-0/0/10        Passed     Down         Fault          1
ge-0/0/11        Passed      UP           OK            0
ge-0/0/12   Not Started     N/A          N/A          N/A
ge-0/0/13   Not Started     N/A          N/A          N/A
ge-0/0/14   Not Started     N/A          N/A          N/A
ge-0/0/15   Not Started     N/A          N/A          N/A
ge-0/0/16   Not Started     N/A          N/A          N/A
ge-0/0/17   Not Started     N/A          N/A          N/A
ge-0/0/18   Not Started     N/A          N/A          N/A
ge-0/0/19        Passed     Down          OK            0
ge-0/0/20   Not Started     N/A          N/A          N/A
ge-0/0/21   Not Started     N/A          N/A          N/A
ge-0/0/22        Passed      UP           OK            0
ge-0/0/23   Not Started     N/A          N/A          N/A
```

Related Documentation

- Troubleshooting Interface Configuration and Cable Faults on page 434
- request diagnostics tdr on page 317
- show diagnostics tdr on page 319

