Network Interfaces Feature Guide for EX2300, EX3400, and EX4300 Switches

Modified: 2017-05-17

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

Network Interfaces Feature Guide for EX2300, EX3400, and EX4300 Switches

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- 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 CLI Explorer.

Documentation Conventions

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
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

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

Convention	Description	Examples
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**.
> (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 of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.
- E-mail—Send your comments to techpubs-comments@juniper.net. 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.

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 22.

This topic describes:

Network Interfaces

Network interfaces connect to the network and carry network traffic. Table 3 on page 19 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.

Table 3: Network Interface Types and Purposes *(continued)*

Type	Purpose
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 20 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.)
Integrated Routing and Bridging (IRB) Interface or Routed VLAN Interface (RVI)	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.

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

Type	Purpose
Virtual Chassis port (VCP) interfaces	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. 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)* . • 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.
Virtual management Ethernet (VME) interface	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.

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 59
- Understanding Layer 3 Subinterfaces on page 106

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 on page 22
- Logical Part of an Interface Name on page 23
- Wildcard Characters in Interface Names on page 24

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 19
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25

CHAPTER 2

Configuring Basic Features on Gigabit Ethernet Interfaces

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

NOTE: This task uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support 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 interface 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 Interface 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 about 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 37 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 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface for trunk port mode:

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

Configuring the Link Settings

EX Series switches include a factory default configuration that enables interfaces with the link settings provided in Table 5 on page 27.

Table 5: Factory Default Configuration Link Settings for EX Series Switches

Ethernet Interface	Autonegotiation	Flow Control	Link Mode	Link Speed
1 gigabit	Enabled	Enabled	Autonegotiation (full duplex or half duplex) For information about EX4300, see the Note below this table.	Autonegotiation (10 Mbps, 100 Mbps, or 1 Gbps)
10 gigabit (using a DAC cable)	Enabled	Enabled	Full duplex	10 Gbps
10 gigabit (using a fiber-optic cable)	Disabled	Enabled	Full duplex	10 Gbps
40 gigabit (using a DAC cable)	Enabled	Enabled	Full duplex	40 Gbps
40 gigabit (using a fiber-optic cable)	Disabled	Enabled	Full duplex	40 Gbps

NOTE: On EX4300 switches, there is no **link-mode** configuration statement. The link-mode setting on an EX4300 switch is handled as follows:

- If the link partner is operating in half duplex, the EX4300 interface goes to half duplex.
- If the link partner is not capable of autonegotiation, the EX4300 interface goes to half duplex.
- If the link partner is capable of autonegotiation and is operating in full duplex, the EX4300 interface also works in full duplex.
- To force an EX4300 interface to stay in full-duplex mode, configure the interface's speed as 10 Mbps or 100 Mbps and also configure the interface with the **no-autonegotiation** statement.

To configure the link mode and speed settings for a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface:

NOTE: On EX4300 switches, there is no link-mode configuration statement. See information earlier in this document regarding how the link mode is set on EX4300 switches.

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

To configure additional link settings for a 1-Gigabit, 10-Gigabit, or 40-Gigabit Ethernet interface:

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

For detailed information about the FPC, PIC, and port numbers used for EX Series switches, see "Understanding Interface Naming Conventions on EX Series Switches" on page 22.

Configurable link settings include:

- **802.3ad**—Specify an aggregated Ethernet bundle. See "Configuring Aggregated Ethernet Links (CLI Procedure)" on page 68.
- **auto-negotiation**—Enable or disable autonegotiation of flow control, link mode, and speed.

NOTE: Starting with Junos OS Releases 14.1X53-D40, 15.1R4, and 17.1R1, half-duplex communication is supported on all built-in network copper ports on EX4300 switches. *Half-duplex* is bidirectional communication; however, signals can flow in only one direction at a time. *Full-duplex* communication means that both ends of the communication can send and receive signals at the same time.

Half-duplex is configured by default on EX4300 switches. If the link partner is set to autonegotiate the link, then the link is autonegotiated to full duplex or half duplex. If the link is not set to autonegotiation, then the EX4300 link defaults to half duplex unless the interface is explicitly configured for full duplex.

To explicitly configure full duplex:

```
[edit]
user@switch# set interfaces interface-name speed 10m-or-100m
[edit]
user@switch# set interfaces interface-name ether-options no-auto-negotiate
```

To verify a half-duplex setting:

```
user@switch> show interfaces interface-name extensive
```

- flow-control—Enable or disable flow control.
- link-mode—Specify full duplex, half duplex, or autonegotiation.

NOTE: On EX4300 switches, there is no link-mode configuration statement. See information earlier in this document regarding how the link mode is set on EX4300 switches.

- loopback—Enable or disable loopback mode.
- speed—Specify 10 Mbps, 100 Mbps, 1 Gbps, 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 EX4300 switches are set to family ethernet-switching by default. You might have to delete this or any other user-configured family setting before changing the setting to family inet or family inet6.

Release History Table

Release	Description
14.1X53-40	Starting with Junos OS Releases 14.1X53-D40, 15.1R4, and 17.1R1, half-duplex communication is supported on all built-in network copper ports on EX4300 switches.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 30
- Monitoring Interface Status and Traffic on page 123
- show interfaces ge- on page 278
- show interfaces xe- on page 313
- Understanding Interface Naming Conventions on EX Series Switches on page 22

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 6 on page 32 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 6: 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. To enable security configuration, select the **Enable Security Configuration** check box. The **forwarding-options dhcp-security-arp-inspection** will be configured.	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. To enable security configuration, select the **Enable Security Configuration** check box. The **forwarding-options dhcp-security groups** and **forwarding-options dhcp-security-arp-inspection** will be configured.	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**.

Table 6: Port Edit Options *(continued)*

Field	Function	Your Action
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 7 on page 36 for more CoS information. To enable security configuration, select the **Enable Security Configuration** check box. The **forwarding-options dhcp-security groups** and **forwarding-options dhcp-security-arp-inspection** will be configured.	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**.
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 7 on page 36 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.

Table 6: Port Edit Options *(continued)*

Field	Function	Your Action
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		
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**.

Table 6: Port Edit Options *(continued)*

Field	Function	Your Action
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 support **Auto-Negotiation 10M-100M** apart from the values mentioned above.
Duplex	Specifies the link mode.	Select one: **automatic**, **half**, or **full**. NOTE: • For EX4300 switches' link-mode setting, see "Configuring Gigabit Ethernet Interfaces (CLI Procedure)" on page 25.
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.
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**.

Table 6: Port Edit Options *(continued)*

Field	Function	Your Action
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 7: 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.
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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Monitoring Interface Status and Traffic on page 123

- EX Series Switches Interfaces Overview on page 19
- *Junos OS CoS for EX Series Switches Overview*
- Understanding Interface Naming Conventions on EX Series Switches on page 22

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 8 on page 37 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 8: 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	
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`*

Table 8: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands
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 interfaceapply-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-namevlan-id vlan id`
Set port family to **ethernet-switching.** Set Port Mode to **access**.	`set interfaces interfaceunit 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 interfacescheduler-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 interfaceunit 0 classifiers dscp juniper-dscp-classifier`

Set CoS Configuration	Refer Table 9 on page 40 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`
Set VLAN on VLANs stanza.	`set vlans vlan namevlan-id vlan-id`

Table 8: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands
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 interfaceunit 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 interfacescheduler-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 interfaceunit 0 classifiers dscp juniper-dscp-classifier`

Set CoS configuration	Refer Table 9 on page 40 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`

Table 8: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands
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`

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 interfacescheduler-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 interfaceunit 0 classifiers dscp juniper-dscp-classifier`

Set CoS configuration	Refer to Table 9 on page 40 for details.

Table 9 on page 40 lists the CLI commands for the recommended CoS settings that are committed when the CoS configuration is set.

Table 9: 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	

Table 9: Recommended CoS Settings for Port Roles *(continued)*

CoS Parameter	CLI Command

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`

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 30
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25

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 Broadband Subscriber Management and Services Library*.

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

Disabling a Physical Interface

Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration.

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.

To disable a physical interface:

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

```
[edit]
user@host# edit interfaces ge-fpc/pic/port
```

2. Include the **disable** statement.

```
[edit interfaces at-fpc/pic/port ]
user@host# set disable
```

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.

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

Verifying the interface configuration:

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

Effect of Disabling Interfaces on T series PICs

The following table describes the effect of using the **set interfaces disable *interface_name*** statement on T series PICs.

Table 10: 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

Related Documentation

- disable on page 163

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 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, perform the following steps:

1. Configure the interface address at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level.

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

     ```
     [edit interfaces ]
     user@host# set interface-name unit logical-unit-number family inet address a.b.c.d/nn
     ```

     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.
     - 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 .
     - 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.
     - 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.

   - To configure an IPv6 address on routers and switches running Junos OS, use the **interface *interface-name* unit *number* family inet6 address *aaaa:bbbb:...:zzzz/nn*** statement at the **[edit interfaces]** hierarchy level.

     ```
     [edit interfaces ]
     user@host# set interface-name unit logical-unit-number family inet6 address
       aaaa:bbbb:...:zzzz/nn
     ```

     NOTE:
     - You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values. The double colon (::) represents all bits set to 0.
     - You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

2. [Optional] Set the broadcast address on the network or subnet .

   [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*],
   user@host# **set** broadcast ***address***

   NOTE: The broadcast 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

3. [Optional] specify the remote address of the connection for the encrypted, PPP-encapsulated, and tunnel interfaces.

   [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]
   user@host# **set destination** ***address***

4. [Optional] For interfaces that carry IP version 6 (IPv6) traffic, configure the host to assign iteslf a unique 64-Bit IP Version 6 interface identifier (EUI-64).

   [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]
   user@host# **set** eui-64

Related Documentation

- *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 Accounting for the Logical Interface

Accounting Profiles Overview

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*.

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*.

Configuring Accounting for the Logical Interface

Before you begin

You must configure a profile to collect error and statistic information for input and output packets on a particular logical interface. An accounting profile specifies what statistics should be collected and written to a log file. For more information on how to configure an accounting-data log file, see the *Configuring Accounting-Data Log Files*.

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular logical interface.

1. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

   ```
   [edit accounting-options interface-profile profile-name]
   user@host# set fieldsfield-name
   ```

2. Each accounting profile logs its statistics to a file in the **/var/log** directory. To configure which file to use, include the **file** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level.

   ```
   [edit accounting-options interface-profile profile-name]
   user@host# set file filename
   ```

   NOTE: You must specify a file statement for the interface profile that has already been configured at the [edit accounting-options] hierarchy level. For more information, see the Configuring Accounting-Data Log Files

3. Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the interval statement at the [**edit accounting-options interface-profile *profile-name***] hierarchy level.

   ```
   [edit accounting-options interface-profile profile-name]
   user@host# set interval minutes
   ```

   NOTE: The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

4. To configure the interfaces on which the accounting needs to be performed, apply the interface profile to a logial interface by including the **accounting-profile** statement at the **[edit interfaces interface-name unit *logical-unit-number*]** hierarchy level.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-numberaccounting-profile profile-name
```

Displaying Accounting Profile for the Logical Interface

Purpose To display the configured accounting profile a particular logical interface at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level:

- interface-name—ge-1/0/1
- Logical unit number—1
- Interface profile —**if_profile**
- File name—**if_stats**
- Interval—15 minutes

Action

- Run the **show** command at the **[edit interfaces ge-1/0/1 unit 1]** hierarchy level.

```
[edit interfaces ge-1/0/1 unit 1]
accounting-profile if_profile;
```

- Run the **show** command at the **[edit accounting-options]** hierarchy level.

```
interface-profile if_profile {
    interval 15;
    file if_stats {
        fields {
            input-bytes;
            output-bytes;
            input-packets;
            output-packets;
            input-errors;
            output-errors;
        }
    }
}
```

Meaning The configured accounting and its associated set options are displayed as expected.

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 207
- *Ethernet Interfaces Overview*
- EX Series Switches Interfaces Overview on page 19
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests help detect duplicate IP addresses. A gratuitous ARP 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. However, 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.

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.

To enable updating of the ARP cache for gratuitous ARPs:

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

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

2. Include the **gratuitous-arp-reply** statement.

   ```
   [edit interfaces interface-name]
   user@host# set gratuitous-arp-reply
   ```

To restore the default behavior, that is, to disable updating of the ARP cache for gratuitous ARP, delete the **gratuitous-arp-reply** statement from the configuration:

```
[edit interfaces interface-name]
user@host# delete gratuitous-arp-reply;
```

By default, the router or switch responds to gratuitous ARP requests. However, on Ethernet interfaces, you can disable responses to gratuitous ARP requests.

To disable responses to gratuitous ARP requests:

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

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

2. Include the **no-gratuitous-arp-request** statement.

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

Related Documentation

- gratuitous-arp-reply on page 179
- no-gratuitous-arp-request on page 218
- *Ethernet Interfaces Overview*
- EX Series Switches Interfaces Overview on page 19
- *Ethernet Interfaces Feature Guide for Routing Devices*

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 178
- *Ethernet Interfaces Overview*
- EX Series Switches Interfaces Overview on page 19
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses

By default, the device responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the device to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

Also, unlike dynamically learned ARP entries, static ARP entries do not age out. You can also configure static ARP entries in a troubleshooting situation or if your device is unable to learn a MAC address dynamically.

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.

To configure static ARP entries:

1. In the configuration mode, at the [**edit**] hierarchy level, configure the router interface on which the ARP table entries for the router is configured.

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

2. Configure the protocol family, the logical unit of the interface, and the interface address of the router interface at the [**edit interfaces** *interface-name*] hierarchy level. While configuring the protocol family, specify **inet** as the protocol family.

   

   NOTE: 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.

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

3. Configure a static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. The IP address specified must be part of the subnet defined in the enclosing **address** statement. The MAC address must be specified as hexadecimal bytes in the following formats: ***nnnn.nnnn.nnnn*** or ***nn:nn:nn:nn:nn:nn***format. For instance, you can use either **0011.2233.4455** or **00:11:22:33:44:55**.

[edit interfaces *interface-name* unit *logical-unit-number* family inet address *interface-address*
user@host# set arp *ip-address* mac *mac-address*

4. Configure another static ARP entry by specifying the IP address and the MAC address that are to be mapped to each other. You can also associate a multicast MAC address with a unicast IP address by including the **multicast-mac** option with the **arp** statement. You can optionally configure the router to respond to ARP requests for the specified IP address by using the **publish** option with the **arp** statement.

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

[edit interfaces *interface-name* unit *logical-unit-number* family inet address *interface-address*
user@host# set arp *ip-address* multicast-mac *mac-address* publish

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

Related Documentation

- arp on page 152
- *Static ARP Table Entries Overview*
- *Management Ethernet Interface Overview*
- EX Series Switches Interfaces Overview on page 19
- *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 219

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 52 for information about how to disable responses to gratuitous ARP requests.

Related Documentation
- proxy-arp on page 224
- *Restricted and Unrestricted Proxy ARP Overview*
- Configuring Gratuitous ARP on page 52
- *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*]

CHAPTER 3

Configuring Aggregated Ethernet Interfaces

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

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 11 on page 60 lists the EX Series switches and the maximum number of interfaces per LAG and the maximum number of LAGs they support. Table 12 on page 61 lists the MX Series routers and the maximum number of interfaces per LAG and the maximum number of LAG groups they support. MX Series routers can support up to 64 LAGs.

Table 11: 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	32
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

Table 12: Maximum Interfaces per LAG and Maximum LAGs per Router

Router	Maximum Interfaces per LAG	Maximum LAG Groups
MX5, MX10, MX40, MX80 and MX104	16	Limited by the interface capacity. 80 on MX104.
MX240, MX480, MX960, MX2010, and MX2020	64 in Junos OS Release 12.3R3	480 in Junos OS Release 9.5R1 1000 in Junos OS Release 14.2R3 1000 in Junos OS Release 16.1R1

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)" on page 82.

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

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 68
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 72
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- *Junos OS Network Interfaces Configuration Guide*

Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic

Juniper Networks EX Series and QFX Series use a hashing algorithm to determine how to forward traffic over a link aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. You can configure some of the fields that are used by the hashing algorithm.

NOTE: **Platform support depends on the Junos OS release in your installation.**

This topic contains the following sections:

Understanding the Hashing Algorithm

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

For ECMP, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The hashing algorithm makes hashing decisions based on values in various packet fields, as well as on some internal values like source port ID and source device ID. The packet fields used by the hashing algorithm varies by the packet's EtherType and, in some instances, by the configuration on the switch. The hashing algorithm recognizes the following EtherTypes:

- IP (IPv4 and IPv6)
- MPLS
- MAC-in-MAC

Traffic that is not recognized as belonging to any of these EtherTypes is hashed based on the Layer 2 header. IP and MPLS traffic are also hashed based on the Layer 2 header when a user configures the hash mode as Layer 2 header.

You can configure some fields that are used by the hashing algorithm to make traffic forwarding decisions. You cannot, however, configure how certain values within a header are used by the hashing algorithm.

Note the following points regarding the hashing algorithm:

- The fields selected for hashing are based on the packet type only. The fields are not based on any other parameters, including forwarding decision (bridged or routed) or egress LAG bundle configuration (Layer 2 or Layer 3).
- The same fields are used for hashing unicast and multicast packets. Unicast and multicast packets are, however, hashed differently.
- The same fields are used by the hashing algorithm to hash ECMP and LAG traffic, but the hashing algorithm hashes ECMP and LAG traffic differently. LAG traffic uses a trunk hash while ECMP uses ECMP hashing. Both LAG and ECMP use the same RTAG7 seed but use different offsets of that 128B seed to avoid polarization. The initial config of

the HASH function to use the trunk and ECMP offset are set at the PFE Init time. The different hashing ensures that traffic is not polarized when a LAG bundle is part of the ECMP next-hop path.

- The same fields are used for hashing regardless of whether the switch is or is not participating in a mixed or non-mixed Virtual Chassis or Virtual Chassis Fabric (VCF).

The fields used for hashing by each EtherType as well as the fields used by the Layer 2 header are discussed in the following sections.

IP (IPv4 and IPv6)

Payload fields in IPv4 and IPv6 packets are used by the hashing algorithm when IPv4 or IPv6 packets need to be placed onto a member link in a LAG bundle or sent to the next-hop device when ECMP is enabled.

The hash mode is set to Layer 2 payload field, by default. IPv4 and IPv6 payload fields are used for hashing when the hash mode is set to Layer 2 payload.

If the hash mode is configured to Layer 2 header, IPv4, IPv6, and MPLS packets are hashed using the Layer 2 header fields. If you want incoming IPv4, IPv6, and MPLS packets hashed by the source MAC address, destination MAC address, or EtherType fields, you must set the hash mode to Layer 2 header.

Table 13 on page 64 displays the IPv4 and IPv6 payload fields that are used by the hashing algorithm, by default.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

Table 13: IPv4 and IPv6 Hashing Fields

Fields	EX4300		QFX5100		QFX5110		QFX5200	
	LAG	ECMP	LAG	ECMP	LAG	ECMP	LAG	ECMP
Source MAC	X	X	X	X	X	X	X	X
Destination MAC	X	X	X	X	X	X	X	X
EtherType	X	X	X	X	X	X	X	X
VLAN ID	X (configurable)	X (configurable)	X (configurable)	X (configurable)	X (configurable)	X (configurable)	X (configurable)	X (configurable)
Source IP or IPv6	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	X (configurable)	X (configurable)

Table 13: IPv4 and IPv6 Hashing Fields *(continued)*

Fields	EX4300		QFX5100		QFX5110		QFX5200	
Destination IP or IPv6	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	X (configurable)	X (configurable)
Protocol (IPv4 only)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	X (configurable)	X (configurable)
Next header (IPv6 only)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	X (configurable)	X (configurable)
Layer 4 Source Port	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	X (configurable)	X (configurable)
Layer 4 Destination Port	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)	X (configurable)	X (configurable)
IPv6 Flow label (IPv6 only)	X	X	X	X	X	X	X	X

MPLS

The hashing algorithm hashes MPLS packets using the source IP, destination IP, MPLS label 0, MPLS label 1, and MPLS label 2 fields. On the QFX5110 and QFX5200 switches, , LSR routers also support ECMP. ECMP uses these fields for hashing on an LSR router:

- Layer 3 VPN: MPLS Labels (top 3 labels), source IP, destination IP, and ingress port ID
- Layer 2 Circuit: MPLS Labels (top 3 labels) and ingress port ID

Table 14 on page 66 displays the MPLS payload fields that are used by the hashing algorithm, by default:

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

The fields used by the hashing algorithm for MPLS packet hashing are not user-configurable.

The source IP and destination IP fields are not always used for hashing. For non-terminated MPLS packets, the payload is checked if the bottom of stack (BoS) flag is seen in the packet. If the payload is IPv4 or IPv6, then the IP source address and IP destination address fields are used for hashing along with the MPLS labels. If the BoS flag is not seen in the packet, only the MPLS labels are used for hashing.

Table 14: MPLS Hashing Fields

Field	EX4300		QFX5100	QFX5110	QFX5200
Source MAC	X		X	X	X
Destination MAC	X		X	X	X
EtherType	X		X	X	X
VLAN ID	X		X	X	X
Source IP	✓		✓	✓	✓
Destination IP	✓		✓	✓	✓
Protocol (for IPv4 packets)	X		X	X	X
Next header (for IPv6 packets)	X		X	X	X
Layer 4 Source Port	X		X	X	X
Layer 4 Destination Port	X		X	X	X
IPv6 Flow lab	X		X	X	X
MPLS label 0	✓		✓	✓	✓
MPLS label 1	✓		✓	✓	✓
MPLS label 2	✓		✓	✓	✓
Ingress Port ID	X	X	X	✓ (LSR and L2Circuit)	✓ (LSR and L2Circuit)

MAC-in-MAC Packet Hashing

Packets using the MAC-in-MAC EtherType are hashed by the hashing algorithm using the Layer 2 payload source MAC, Layer 2 payload destination MAC, and Layer 2 payload EtherType fields. See Table 15 on page 67.

Hashing using the fields in the MAC-in-MAC EtherType packet is first supported on EX4300 switches in Release 13.2X51-D20. Hashing using the fields in the MAC-in-MAC EtherType is not supported on earlier releases.

The fields used by the hashing algorithm for MAC-in-MAC hashing are not user-configurable.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.

Table 15: MAC-in-MAC Hashing Fields

Field	EX4300	QFX5100	QFX5110	QFX5200
Layer 2 Payload Source MAC	✓	✓	✓	✓
Layer 2 Payload Destination MAC	✓	✓	✓	✓
Layer 2 Payload EtherType	✓	✓	✓	✓
Layer 2 Payload Outer VLAN	X	X	X	X

Layer 2 Header Hashing

Layer 2 header fields are used by the hashing algorithm when a packet's EtherType is not recognized as IP (IPv4 or IPv6), MPLS, or MAC-in-MAC. The Layer 2 header fields are also used for hashing IPv4, IPv6, and MPLS traffic instead of the payload fields when the hash mode is set to Layer 2 header.

- ✓—Field is used by the hashing algorithm, by default.
- X—Field is not used by the hashing algorithm, by default.
- (configurable)—Field can be configured to be used or not used by the hashing algorithm.

Table 16: Layer 2 Header Hashing Fields

Field	EX4300	QFX5100	QFX5110	QFX5200
Source MAC	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
Destination MAC	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
EtherType	✓ (configurable)	✓ (configurable)	✓ (configurable)	✓ (configurable)
VLAN ID	X (configurable)	X (configurable)	✓ (configurable)	✓ (configurable)

Related Documentation

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 82

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 72.

Related Documentation

- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 69
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 72
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Verifying the Status of a LAG Interface on page 127
- Understanding Aggregated Ethernet Interfaces and LACP on page 59

Configuring Aggregated Ethernet Interfaces (J-Web Procedure)

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

For J-Web Application package Release 14.1X53-A2, on EX4600 switches the maximum number of link aggregation group (LAG) devices supported is 1000.

Use the link aggregation feature to aggregate one or more Ethernet interfaces to form a virtual link or 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 when failure occurs, and increases availability. You can use the J-Web interface to configure LAGs, 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 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 (J-Web Procedure)* for details about all commit options.

2. Select one of the following:

- **Add**—Creates a LAG. Enter information as specified in Table 17 on page 70.
- **Edit**—Modifies a selected LAG.
  - **Aggregation**—Modifies settings for the selected LAG. Enter information as specified in Table 17 on page 70.
  - **VLAN**—Specifies VLAN options for the selected LAG. Enter information as specified in Table 18 on page 71.
  - **IP Option**—Specifies IP options for the selected LAG. Enter information as specified in Table 19 on page 71.
- **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 17: 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 Link Aggregation Control Protocol (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 18: 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 19: 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**.

Release History Table

Release	Description
14.1X53-A2	For J-Web Application package Release 14.1X53-A2, on EX4600 switches the maximum number of link aggregation group (LAG) devices supported is 1000.

Related Documentation

- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Verifying the Status of a LAG Interface on page 127
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 72
- Understanding Aggregated Ethernet Interfaces and LACP on page 59

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 73.

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 68

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 68
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 69
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Verifying the Status of a LAG Interface on page 127
- Understanding Aggregated Ethernet Interfaces and LACP on page 59

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 68.

- Configured LACP for the interface. See "Configuring Aggregated Ethernet LACP (CLI Procedure)" on page 72.

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 on page 75
- Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level on page 75
- Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface on page 76

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 59

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 78
- Configuring Primary and Backup Links for Link Aggregated Ethernet Interfaces on page 78
- Reverting Traffic to a Primary Link When Traffic is Passing Through a Backup Link on page 79
- Disabling Link Protection for Aggregated Ethernet Interfaces on page 79

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, there are exceptions.

Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers. For example, these mixes are supported:

- 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)

Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.

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, TX Matrix Plus, and PTX routers, you need to configure the **mixed** option for the [**edit interfaces ae*x* 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, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches 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.

Release History Table

Release	Description
14.2	Starting with Junos OS Release 14.2, aggregated Ethernet supports mixed link speeds on PTX Series Packet Transport Routers.
13.2	Starting with Junos OS Release 13.2, aggregated Ethernet supports mixed rates and mixed modes on T640, T1600, T4000, and TX Matrix Plus routers.
13.2	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.

Related Documentation

- *aggregated-ether-options*
- *Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles*
- *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**.

On MX Series routers, when Link Aggregation Control Protocol (LACP) is enabled on a link aggregation group (LAG) interface along with minimum links configuration, the bundle is considered to be up when the following two conditions are met:

- The specified minimum number of links are up.
- The links are in *collecting distributing* state—that is, collecting and distributing states are merged together to form a combined state (coupled control) for the aggregated port. 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.

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*
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure)

Juniper Networks EX Series and QFX Series switches use a hashing algorithm to determine how to forward traffic over a Link Aggregation group (LAG) bundle or to the next-hop device when equal-cost multipath (ECMP) is enabled.

The hashing algorithm makes hashing decisions based on values in various packet fields.. You can configure some of the fields that are used by the hashing algorithm.

Configuring the fields used by the hashing algorithm is useful in scenarios where most of the traffic entering the bundle is similar and the traffic needs to be managed in the LAG bundle. For instance, if the only difference in the IP packets for all incoming traffic is the source and destination IP address, you can tune the hashing algorithm to make hashing decisions more efficiently by configuring the algorithm to make hashing decisions using only those fields.

NOTE: **Configuring the hash mode is not supported on QFX10002 and QFX10008 switches.**

- Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing on page 83
- Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing on page 84
- Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing on page 84

Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing

To configure the hashing algorithm to use fields in the Layer 2 header for hashing:

1. Configure the hash mode to Layer 2 header:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set hash-mode layer2-header
   ```

   The default hash mode is Layer 2 payload. Therefore, this step must be performed if you have not previously configured the hash mode.

2. Configure the fields in the Layer 2 header that the hashing algorithm uses for hashing:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set layer2 {no-destination-mac-address | no-ether-type |
   no-source-mac-address | vlan-id}
   ```

   By default, the hashing algorithm uses the values in the destination MAC address, Ethertype, and source MAC address fields in the header to hash traffic on the LAG. You can configure the hashing algorithm to not use the values in these fields by configuring **no-destination-mac-address**, **no-ether-type**, or **no-source-mac-address**.

   You can also configure the hashing algorithm to include the VLAN ID field in the header by configuring the **vlan-id** option.

   If you want the hashing algorithm to not use the Ethertype field for hashing:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set layer2 no-ether-type
   ```

Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing

To configure the hashing algorithm to use fields in the IP payload for hashing:

1. Configure the hash mode to Layer 2 payload:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set hash-mode layer2-payload
   ```

   The IP payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IP payload that the hashing algorithm uses for hashing:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set inet {no-ipv4-destination-address | no-ipv4-source-address |
   no-l4-destination-port | no-l4-source-port | no-protocol | vlan-id}
   ```

   For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and protocol fields and instead hash traffic based only on the IPv4 source and destination addresses:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set inet no-l4-destination-port no-l4-source-port no-protocol
   ```

Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing

To configure the hashing algorithm to use fields in the IPv6 payload for hashing:

1. Configure the hash mode to Layer 2 payload:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set hash-mode layer2-payload
   ```

   The IPv6 payload is not checked by the hashing algorithm unless the hash mode is set to Layer 2 payload. The default hash mode is Layer 2 payload.

2. Configure the fields in the IPv6 payload that the hashing algorithm uses for hashing:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set inet6 {no-ipv6-destination-address | no-ipv6-source-address |
   no-l4-destination-port | no-l4-source-port | no-next-header | vlan-id}
   ```

   For instance, if you want the hashing algorithm to ignore the Layer 4 destination port, Layer 4 source port, and the Next Header fields and instead hash traffic based only on the IPv6 source and IPv6 destination address fields only:

   ```
   [edit forwarding-options enhanced-hash-key]
   user@switch# set inet6 no-l4-destination-port no-l4-source-port no-next-header
   ```

Related Documentation

- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 62
- *Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)*
- Understanding Aggregated Ethernet Interfaces and LACP on page 59

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 aex]
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 105.

Related Documentation

- *vlan-id*
- vlan-tagging on page 238

CHAPTER 4

Configuring Energy Efficient Interfaces

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 87

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 88
- Disabling EEE on a Base-T Copper Ethernet Port on page 88

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 127
- Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 87

CHAPTER 5

Configuring Interface Ranges

Understanding Interface Ranges on EX Series Switches

NOTE: This concept uses Junos OS for EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support 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, 40-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:

- **forwarding-options analyzer *name* input egress interface**
- **forwarding-options analyzer *name* input ingress interface**

- poe interface
- protocols dot1x authenticator interface
- protocols igmp interface
- protocols isis interface
- protocols layer2-control bpdu-block interface
- protocols link-management peer *name* lmp-control-channel
- protocols link-management te-link *name* interface
- protocols lldp interface
- protocols lldp-med interface
- protocols mstp interface
- protocols oam ethernet link-fault-management interface
- protocols ospf area *area-id* interface
- protocols pim interface
- protocols router-advertisement interface
- protocols router-discovery interface
- protocols rsvp interface
- protocols sflow interfaces
- protocols vstp vlan *vlan-id* interface
- switch-options redundant-trunk-group *group-name* interface
- switch-options voip interface

Related Documentation

- Configuring Interface Ranges on page 90
- EX Series Switches Interfaces Overview on page 19
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 106
- interface-range on page 193

Configuring Interface Ranges

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

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 on Switches

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.

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 must 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*.

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

- **forwarding-options analyzer** *name* **input egress interface**
- **forwarding-options analyzer** *name* **input ingress interface**
- **poe interface**
- **protocols dot1x authenticator interface**
- **protocols igmp interface**
- **protocols isis interface**
- **protocols layer2-control bpdu-block interface**
- **protocols link-management peer** *name* **lmp-control-channel**
- **protocols link-management te-link** *name* **interface**
- **protocols lldp interface**
- **protocols lldp-med interface**
- **protocols mstp interface**
- **protocols oam ethernet link-fault-management interface**
- **protocols ospf area** *area-id* **interface**
- **protocols pim interface**
- **protocols router-advertisement interface**
- **protocols router-discovery interface**
- **protocols rsvp interface**
- **protocols sflow interfaces**
- **protocols vstp vlan** *vlan-id* **interface**
- **switch-options redundant-trunk-group group-name interface**
- **switch-options voip 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.

CHAPTER 6

Configuring IP Directed Broadcast

Understanding IP Directed Broadcast

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 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

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 a Switch*
- Configuring IP Directed Broadcast (CLI Procedure) on page 103

Configuring IP Directed Broadcast (CLI Procedure)

NOTE: This task uses Junos OS with support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support 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 a 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 an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) for the subnet that will be enabled for IP direct broadcast. See *Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)*, *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 irb.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) irb.1
```

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

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

Related Documentation

- *Example: Configuring IP Directed Broadcast on a Switch*
- Understanding IP Directed Broadcast on page 101

CHAPTER 7

Configuring Layer 3 Subinterfaces

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*
- *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 Access Mode on a Logical Interface*
- *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 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 19
- *Junos OS Ethernet Interfaces Configuration Guide*

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.

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*
- Verifying That Layer 3 Subinterfaces Are Working on page 131
- Understanding Layer 3 Subinterfaces on page 106

CHAPTER 8

Configuring Local Link Bias

Understanding Local Link Bias

NOTE: The QFX5200 switches do not support Virtual Chassis or Virtual Chassis ports.

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 traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is enabled is illustrated in Figure 1 on page 109.

Figure 1: Egress Traffic Flow with Local Link Bias

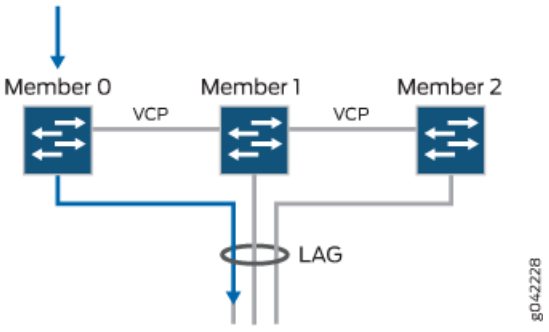

When local link bias is disabled, egress traffic exiting a Virtual Chassis or VCF on a LAG bundle can be forwarded out of any member link in the LAG bundle. Traffic forwarding decisions are made by an internal algorithm that attempts to load-balance traffic between the member links in the bundle. VCP 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 traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is disabled is illustrated in Figure 2 on page 110.

Figure 2: Egress Traffic Flow without Local Link Bias

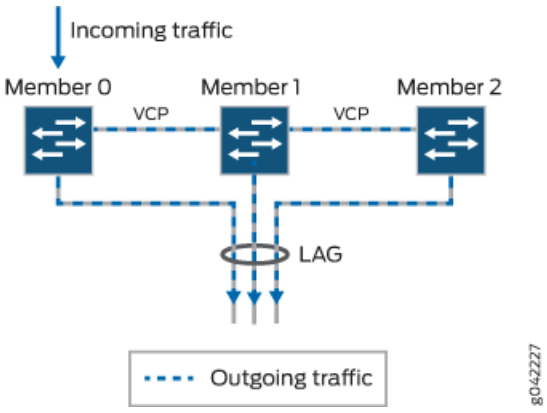

Starting in Junos OS Release 14.1X53-D25, local link bias can be enabled globally for all LAG bundles in a Virtual Chassis or VCF, or individually per LAG bundle in a Virtual Chassis. In prior Junos OS releases, local link bias could be enabled individually per LAG bundle only.

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 a 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.

Release History Table

Release	Description
14.1X53-D25	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.

Related Documentation

- Configuring Local Link Bias (CLI Procedure) on page 111

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 a LAG 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.

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**.

To enable local link bias on a LAG bundle:

```
[edit]
user@switch# set interface aex aggregated-ether-options local-bias
```

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
```

Related Documentation

- Understanding Local Link Bias on page 109

CHAPTER 9

Configuring Unicast RPF

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, by default 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. You can also enable loose mode, which means that the system checks to see if the packet has a source address with a corresponding prefix in the routing table but does not check whether the receiving interface is the best return path to the packet's unicast source address.

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 117.

NOTE: Platform support depends on the Junos OS release in your installation.

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 115.)

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 3 on page 115. 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 3: Symmetrically Routed Interfaces

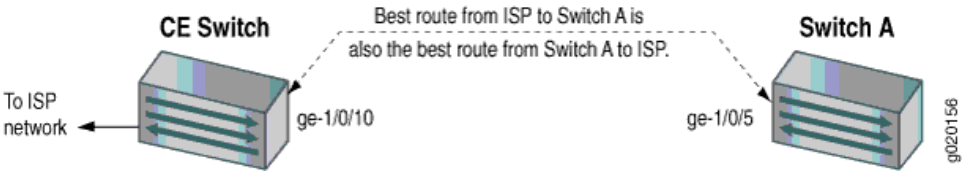

Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

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 4 on page 116. 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 4: Asymmetrically Routed Interfaces

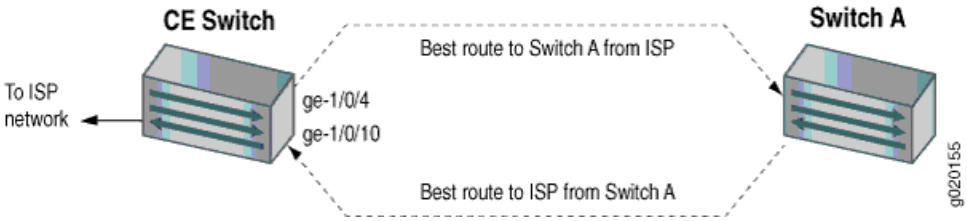

NOTE: Do not enable unicast RPF on EX3200, EX4200, and EX4300 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of these switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of 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*
- Configuring Unicast RPF (CLI Procedure) on page 117
- Disabling Unicast RPF (CLI Procedure) on page 119

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. When you enable unicast RPF, by default 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. You can also enable loose mode, which means that the system checks to see if the packet has a source address with a corresponding prefix in the routing table but does not check whether the receiving interface is the best return path to the packet's unicast source address.

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 interface-name unit 0 family inet rpf-check
```

To enable unicast RPF loose mode, enter:

```
[edit interfaces]
user@switch# set interface-name unit 0 family inet rpf-check mode loose
```

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*
- Verifying Unicast RPF Status on page 132
- Disabling Unicast RPF (CLI Procedure) on page 119
- Troubleshooting Unicast RPF on page 137
- Understanding Unicast RPF on page 113

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*
- Verifying Unicast RPF Status on page 132

- Configuring Unicast RPF (CLI Procedure) on page 117
- Understanding Unicast RPF on page 113

PART 1

Troubleshooting Information

CHAPTER 10

Monitoring and Troubleshooting Interfaces

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 30
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25

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 125
- *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*.

Related Documentation

- *Tracing Interface Operations Overview*
- *Tracing Operations of an Individual Router Interface*
- traceoptions on page 231

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 68
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 69
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

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 87
- Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces on page 87

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

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

1. Verifying the LACP Setup on page 129
2. Verifying That LACP Packets Are Being Exchanged on page 130

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 72
- *Verifying the Status of a LAG Interface*
- Verifying the Status of a LAG Interface on page 127

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     10.1.1.1/24
ge-0/0/0.1              up    up   inet     10.1.1.2/24
ge-0/0/0.2              up    up   inet     10.1.1.3/24
ge-0/0/0.3              up    up   inet     10.1.1.4/24
ge-0/0/0.4              up    up   inet     10.1.1.5/24
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 10.1.1.1 (10.1.1.1): 56 data bytes
64 bytes from 10.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 10.1.1.1 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 106
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*

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 example below 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 and EX4200 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs) 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 313
- *Example: Configuring Unicast RPF on an EX Series Switch*
- *Configuring Unicast RPF on ACX Series Routers*
- Configuring Unicast RPF (CLI Procedure) on page 117
- Disabling Unicast RPF (CLI Procedure) on page 119
- Troubleshooting Unicast RPF on page 137

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 a Switch*.

Related Documentation

- *Configuring IP Directed Broadcast (CLI Procedure)*
- Configuring IP Directed Broadcast (CLI Procedure) on page 103
- *Example: Configuring IP Directed Broadcast on a Switch*

Troubleshooting an Aggregated Ethernet Interface

Troubleshooting issues for aggregated Ethernet interfaces:

- Show Interfaces Command Shows the LAG is Down on page 134
- Logical Interface Statistics Do Not Reflect All Traffic on page 135
- IPv6 Interface Traffic Statistics Are Not Supported on page 135
- SNMP Counters ifHCInBroadcastPkts and ifInBroadcastPkts Are Always 0 on page 135

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 127
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

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 136

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 123
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 30
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- *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 137

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 132
- Understanding Unicast RPF on page 113

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 136
- request diagnostics tdr on page 252
- show diagnostics tdr on page 254

PART 2

Configuration Statements and Operational Commands

CHAPTER 11

Configuration Statements

- vlan (802.1Q Tagging) on page 236
- vlan-id (VLAN Tagging and Layer 3 Subinterfaces) on page 237
- vlan-tagging on page 238

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 aggregration group (LAG).

Options

- **aex**—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*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 72
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73

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.
Statement introduced in Junos OS Release 15.1F6 for PTX Series routers with third-generation FPCs installed.

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*
- Applying an Accounting Profile to the Logical Interface on page 49

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;
  virtual-gateway-address
  (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.

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, and 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 192.168.1.1/8;
      }
    }
  }
  xe-0/0/1 {
    unit 0 {
     family inet {
       address 192.168.1.1/8;
     }
  }
}
```

For more information on configuring the same address on multiple interfaces, see "Configuring the Interface Address" on page 46.

The remaining statements are explained separately. See CLI Explorer.

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*
- *family*
- *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. See CLI Explorer.

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*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- Understanding Aggregated Ethernet Interfaces and LACP on page 59
- *Junos OS Ethernet Interfaces Configuration Guide*

aggregated-ether-options

Syntax

```
aggregated-ether-options {
  ethernet-switch-profile {
    tag-protocol-id;
  }
  (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;
  local-bias;
  logical-interface-fpc-redundancy;
  (loopback | no-loopback);
  mc-ae {
    chassis-id chassis-id;
    events {
      iccp-peer-down {
        force-icl-down;
        prefer-status-control-active;
      }
    }
    init-delay-time seconds;
    mc-ae-id mc-ae-id;
    mode (active-active | active-standby);
    redundancy-group group-id;
    revert-time revert-time;
    status-control (active | standby);
    switchover-mode (non-revertive |revertive);
  }
  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.
Statement introduced in Junos OS Release 12.3R2.

Description Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.

The remaining statements are explained separately. See CLI Explorer.

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*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 72
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- *Configuring Q-in-Q Tunneling (CLI Procedure)*
- *Junos OS Ethernet Interfaces Configuration Guide*

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*]

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

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.

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.

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, **0000.5e00.5355** or **00:00:5e:00:53:55**.

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, **0000.5e00.5355** or **00:00:5e:00:53:55**.

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

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 For Mapping IP Addresses to MAC Addresses on page 55

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 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, 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 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 (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25

backup-liveness-detection

Syntax
```
backup-liveness-detection {
  backup-peer-ip ipv4-address;
}
```

Hierarchy Level [edit protocols iccp peer]

Release Information Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 13.2R1 for EX Series switches.

Description Determine whether a peer is up or down by exchanging keepalive messages over the management link between the two Inter-Chassis Control Protocol (ICCP) peers.

When an ICCP connection is operationally down, the status of the peers hosting a multichassis link aggregation group (MC-LAG) is detected by sending liveness detection requests to each other. Peers must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, the liveness detection check fails, and a failure action is implemented. Backup liveness detection must be configured on both peers hosting the MC-LAG.

The remaining statement is explained separately. See CLI Explorer.

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

Related Documentation
- *Configuring Multichassis Link Aggregation on MX Series Routers*

backup-peer-ip

Syntax backup-peer-ip *ipv4-address*;

Hierarchy Level [edit protocols iccp peer backup-liveness-detection]

Release Information Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 13.2R1 for EX Series switches.

Description Specify the IP address of the peer being used as a backup peer in the Bidirectional Forwarding Detection (BFD) configuration.

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

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 the 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 48

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 (0.0.0.0), nor can you specify a broadcast address (255.255.255.255). For example, in the statement **set interface ge-0/0/0 unit 0 family inet address 10.1.1.0/24**, the subnet address **10.1.1.0** has the host address of 0. Hence, you cannot configure this address. Similarly, for the subnet, you cannot use the broadcast address **10.1.1.255/24**.

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 46

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. See CLI Explorer.

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 68
- *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 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 42
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- *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*

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*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- *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 42
- Disabling a Logical Interface on page 44

enhanced-hash-key

List of Syntax

Syntax (EX Series and QFX5100 Switches) on page 165
Syntax (QFX10002 and QFX10008 Switches) on page 165

Syntax (EX Series and QFX5100 Switches)

```
enhanced-hash-key {
  ecmp-resilient-hash;
  fabric-load-balance {
    flowlet {
      inactivity-interval interval;
    }
    per-packet;
  }
  hash-mode {
    layer2-header;
    layer2-payload;
  }
  inet {
    no-ipv4-destination-address;
    no-ipv4-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-protocol;
    vlan-id;
  }
  inet6 {
    no-ipv6-destination-address;
    no-ipv6-source-address;
    no-l4-destination-port;
    no-l4-source-port;
    no-next-header;
    vlan-id;
  }
  layer2 {
    no-destination-mac-address;
    no-ether-type;
    no-source-mac-address;
    vlan-id;
  }
}
```

Syntax (QFX10002 and QFX10008 Switches)

```
enhanced-hash-key {
  hash-seed seed-value;
  inet {
    no-ipv4-destination-address;
    no-ipv4-source-address;
    no-l4-destination-port;
    no-l4-source-port;
  }
  inet6 {
    ipv6-flow-label;
    no-ipv6-destination-address;
    no-ipv6-source-address;
```

```
    no-l4-destination-port;
    no-l4-source-port;
  }
  layer2 {
    destination-mac-address
    inner-vlan-id;
    no-ether-type;
    no-vlan-id;
    source-mac-address;
  }
  no-mpls;
  gre {
    key;
    protocol;
  }
  vxlan-vnid
  }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.
The **fabric-load-balance** statement introduced in Junos OS Release 14.1X53-D10.
The **hash-seed** statement introduced in Junos OS Release 15.1X53-D30.

Description Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

On QFX10002 and QFX10008 switches, you can configure the hash seed for load balancing.

By default, the QFX10002 and QFX10008 switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the **hash-seed** statement at the [**edit forwarding-options enhanced-hash-key**] hierarchy level. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system will by generate a hash seed value based on the system MAC address.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

Related Documentation

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 82
- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 62

ether-options

Syntax

```
ether-options {
  802.3ad {
    aex;
    (backup | primary);
    lacp {
      force-up;
      port-priority
    }
  }
  (auto-negotiation | no-auto-negotiation);
  ethernet-switch-profile {
    tag-protocol-id;
  }
  (flow-control | no-flow-control);
  ieee-802-3az-eee;
  link-mode mode;
  (loopback | no-loopback);
  speed (speed | auto-negotiation);
}
```

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.
Statement introduced in Junos OS Release 12.3R2.

Description

Configure Ethernet properties for a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface.

The remaining statements are explained separately. See CLI Explorer.

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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 30
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- *Configuring Q-in-Q Tunneling (CLI Procedure)*
- *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. See CLI Explorer.

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 68
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- *Junos OS Ethernet Interfaces Configuration Guide*

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 46

family

Syntax

family ccc

```
family ccc;
  filter {
    group group-number;
    input filter-name;
    input-list [filter-names];
    output filter-name;
    output-list [filter-names];
  }
  policer {
    input policer-name;
    output policer-name;
  }
}
```

family ethernet-switching

```
family ethernet-switching {
  filter {
    input filter-name;
    output filter-name;
  }
  interface-mode (access | trunk);
  recovery-timeout seconds;
  storm-control profile-name;
  vlan {
    members (vlan-name |[-vlan-names] | all);
  }
}
```

family inet

```
family inet {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address ipv4-address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    preferred;
    primary;
    vrrp-group group-number {
      (accept-data | no-accept-data);
      advertise–interval seconds;
      advertisements-threshold number;
      authentication-key key;
```

```
            authentication-type authentication;
            fast-interval milliseconds;
            (preempt | no-preempt) {
              hold-time seconds;
            }
            priority number;
            track {
              interface interface-name {
                priority-cost number;
              }
              priority-hold-time seconds;
              route ip-address/mask routing-instance instance-name priority-cost cost;
            }
            virtual-address [addresses];
            vrrp-inherit-from {
              active-group group-number;
              active-interface interface-name;
            }
          }
        }
        filter {
          input filter-name;
          output filter-name;
        }
        mtu bytes;
        no-neighbor-learn;
        no-redirects;
        primary;
        rpf-check {
          fail-filter filter-name;
          mode {
            loose;
          }
        }
      }
```

family inet6

```
family inet6 {
  accounting {
    destination-class-usage;
    source-class-usage {
      input;
      output;
    }
  }
  address address {
    eui-64;
    ndp ip-address (mac | multicast-mac) mac-address <publish>;
    preferred;
    primary;
    vrrp-inet6-group group-id {
      accept-data | no-accept-data;
      advertisements-threshold number;
      authentication-key key;
      authentication-type authentication;
      fast-interval milliseconds;
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      track {
        interface interface-name {
          priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
      }
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6–address;
      vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
      }
    }
  }
  (dad-disable | no-dad-disable);
  filter {
    input filter-name;
    output filter-name;
  }
  mtu bytes;
  nd6-stale-time time;
  no-neighbor-learn;
  no-redirects;
  policer {
    input policer-name;
    output policer-name;
  }
  rpf-check {
    fail-filter filter-name;
    mode {
      loose;
```

```
    }
  }
}
```

family iso

```
family iso {
  address interface-address;
  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** 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 EX4300 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.

Options See Table 20 on page 175 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 (**ae0**)
- 40-Gigabit Ethernet (**et**)
- Gigabit Ethernet (**ge**)
- Interface-range configuration (**interface-range**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Integrated Routing and Bridging (IRB) interfaces (IRB) (**irb**)
- 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 20: Protocol Families and Supported Interface Types

Family	Description	Supported Interface Types							
		ae0	et	ge	irb	lo0	me0	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	✓	✓	✓	✓	✓	✓	✓	✓

The remaining statements are explained separately. See CLI Explorer.

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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- *Configuring Integrated Routing and Bridging 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.

Description

NOTE: On EX Series switches, the **group**, **input-list**, **output-filter** statements are not supported under the [**edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet**], [**edit interfaces** *interface-name* **unit** *logical-unit-number* **family inet6**], and [**edit interfaces** *interface-name* **unit** *logical-unit-number* **family mpls**] hierarchies.

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.

Options

group ***filter-group-number***—Define an interface to be part of a filter group. The default filter group number is 0.

Range: 0 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. See CLI Explorer.

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 Administration Library*

- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Firewall Filters (CLI Procedure)*
- *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 53
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25

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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 30
- Understanding Aggregated Ethernet Interfaces and LACP on page 59
- *Junos OS Ethernet Interfaces Configuration Guide*

gratuitous-arp-reply

Syntax (gratuitous-arp-reply | no-gratuitous-arp-reply);

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 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 52
- **no-gratuitous-arp-request on page 218**

hash-mode

Syntax

```
hash-mode {
  layer2-header;
  layer2-payload;
}
```

Hierarchy Level [edit forwarding-options enhanced-hash-key]

Release Information Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.
Statement is not supported on QFX10002 and QFX 10008 switches.

Description Select the mode for the hashing algorithm.

The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.

For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.

The hash mode that is set using this statement determines which fields are inspected by the hashing algorithm. You must set the hash mode to **layer2-payload** if you want the hashing algorithm to inspect fields in the Layer 2 payload when making hashing decisions. You must set the hash mode to **layer2-header** if you want the hashing algorithm to inspect fields in the Layer 2 header when making hashing decisions.

If the hash mode is set to **layer2-payload**, you can set the fields used by the hashing algorithm to hash IPv4 traffic using the **set forwarding-options enhanced-hash-key inet** statement. You can set the fields used by the hashing algorithm to hash IPv6 traffic using the **set forwarding-options enhanced-hash-key inet6** statement.

If the hash mode is set to **layer2-header**, you can set the fields that the hashing algorithm inspects in the Layer 2 header using the **set forwarding-options enhanced-hash-key layer2** statement.

Default layer2-payload

Options

layer-2-payload—Set the hashing algorithm to use fields in the Layer 2 payload to make hashing decisions.

layer-2-header—Set the hashing algorithm to use fields in the Layer 2 header to make hashing decisions.

Required Privilege Level

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

Related Documentation

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 82
- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 62
- enhanced-hash-key on page 165
- inet on page 186
- inet6 on page 188
- layer2 on page 198

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.
Statement introduced in Junos OS Release 12.1 for the SRX Series.

Description Specify the **hold-time** value to use to damp shorter interface transitions milliseconds. The hold timer enables interface damping by not advertising interface transitions until the hold timer duration has passed. When a hold-down timer is configured 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 a hold-up timer is configured and 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.

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.

NOTE: On MX Series routers with MPC3E and MPC4E, we recommend that you do not configure the hold-down timer to be less than 1 second. On MX Series routers with MPC5EQ-100G10G (MPC5EQ) or MPC6E (MX2K-MPC6E) with 100-Gigabit Ethernet MIC with CFP2 OTN interfaces, we recommend that you do not configure the hold-down timer to be less than 3 seconds.

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

iccp

Syntax

```
iccp {
  authentication-key string;
  local-ip-addr local-ip-addr;
  peer ip-address{
    authentication-key string;
    backup-liveness-detection {
      backup-peer-ip ip-address;
    }
    liveness-detection {
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      multiplier number;
      no-adaptation;
      transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
      }
      version (1 | automatic);
    }
    local-ip-addr ipv4-address;
    session-establishment-hold-time seconds;
  }
  session-establishment-hold-time seconds;
  traceoptions {
    file <filename> <files number> <match regular-expression> <microsecond-stamp>
      <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit protocols]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure Inter-Chassis Control Protocol (ICCP) between the multichassis link aggregation group (MC-LAG) peers. ICCP replicates forwarding information, validates configurations, and propagates the operational state of the MC-LAG members.

NOTE: Backup liveness detection is not supported on MX Series routers.

The remaining statements are explained separately. See CLI Explorer.

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

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 87

inet (enhanced-hash-key)

Syntax (EX Series and QFX5100 Switch)

```
inet {
  no-ipv4-destination-address;
  no-ipv4-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-protocol;
  vlan-id;
}
```

Syntax (QFX10002 and QFX10008 Switches)

```
inet {
  no-ipv4-destination-address;
  no-ipv4-source-address;
  no-l4-destination-port;
  no-l4-source-port;
}
```

Hierarchy Level [edit forwarding-options enhanced-hash-key]

Release Information Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.
Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and QFX10008 Switches.

Description Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.

When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the **inet** statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to **layer2-payload**. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the **set forwarding-options enhanced-hash-key hash-mode layer2-payload** statement.

Default The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:

- IP destination address
- IP source address
- Layer 4 destination port

- Layer 4 source port
- Protocol

Options

no-ipv4-destination-address—Exclude the IPv4 destination address field from the hashing algorithm.

no-ipv4-source-address—Exclude the IPv4 source address field from the hashing algorithm.

no-l4-destination-port—Exclude the Layer 4 destination port field from the hashing algorithm.

no-l4-source-port—Exclude the Layer 4 source port field from the hashing algorithm.

no-protocol—Exclude the protocol field from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

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

Related Documentation

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 82
- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 62
- *Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)*
- *hash-seed*
- enhanced-hash-key on page 165
- hash-mode on page 180
- inet6 on page 188

inet6 (enhanced-hash-key)

List of Syntax

Syntax (EX Series and QFX5100 Switch) on page 188
Syntax (QFX10002 and QFX10008 Switches) on page 188

Syntax (EX Series and QFX5100 Switch)

```
inet6 {
  no-ipv6-destination-address;
  no-ipv6-source-address;
  no-l4-destination-port;
  no-l4-source-port;
  no-next-header;
  vlan-id;
}
```

Syntax (QFX10002 and QFX10008 Switches)

```
inet6 {
  ipv6-flow-label;
  no-ipv6-destination-address;
  no-ipv6-source-address;
  no-l4-destination-port;
  no-l4-source-port;
}
```

Hierarchy Level

[edit forwarding-options enhanced-hash-key]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 on EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 on QFX Series devices.
Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and QFX 10008 switches.

Description

Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.

When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.

The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the **set forwarding-options enhanced-hash-key hash-mode layer2-payload** statement.

Default

The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:

- IP destination address

- IP source address
- Layer 4 destination port
- Layer 4 source port
- Next header

Options

no-ipv6-destination-address—Exclude the IPv6 destination address field from the hashing algorithm.

no-ipv6-source-address—Exclude the IPv6 source address field from the hashing algorithm.

no-l4-destination-port—Exclude the Layer 4 destination port field from the hashing algorithm.

no-l4-source-port—Exclude the Layer 4 source port field from the hashing algorithm.

no-next-header—Exclude the Next Header field from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

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

Related Documentation

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 82
- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 62
- *Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)*
- *hash-seed*
- enhanced-hash-key on page 165
- hash-mode on page 180
- inet on page 186

interface (Multichassis Protection)

Syntax interface *interface-name*;

Hierarchy Level [edit multi-chassis multi-chassis-protection peer]

Release Information Statement introduced in Junos OS Release 9.6 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the name of the interface that is being used as an interchassis link-protection link (ICL-PL). The two switches hosting a multichassis link aggregation group (MC-LAG) use this link to pass Inter-Chassis Control Protocol (ICCP) and data traffic.

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

interface-mode

Syntax interface-mode (access | trunk <inter-switch-link>);

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

Release Information Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for the QFX Series.
Statement introduced in Junos OS Release 15.1.
inter-switch-link option introduced in Junos OS Release 14.2 for MX240, MX480, and MX960 routers in enhanced LAN mode.

Description

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

(QFX3500 and QFX3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.

NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure **interface-mode** and **irb** for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see *Configuring a Trunk Interface on a Bridge Network*.

Options

access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the **vlan-id** statement.

trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the **vlan-id** or **vlan-id-list** statement.

trunk inter-switch-link—For a private VLAN, configure the InterSwitch Link protocol (ISL) on a trunk port of the primary VLAN in order to connect the switches composing the

PVLAN to each other. You do not need to configure an ISL when a PVLAN is configured on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

Required Privilege Level

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

Related Documentation

- *Configuring Access Mode on a Logical Interface*
- *Configuring a Logical Interface for Trunk Mode*
- *Example: Connecting Access Switches to a Distribution Switch*
- *Tunnel Services Overview*
- *Tunnel Interface Configuration on MX Series Routers Overview*

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. See CLI Explorer.

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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- *Understanding Interface Ranges on EX Series Switches*
- Understanding Interface Ranges on EX Series Switches on page 89
- EX Series Switches Interfaces Overview on page 19
- *Junos OS Interfaces Fundamentals 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.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.
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. See CLI Explorer.

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*

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*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 72
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- Understanding Aggregated Ethernet Interfaces and LACP on page 59
- *Junos OS Ethernet Interfaces Configuration Guide*

layer2 (enhanced-hash-key)

List of Syntax

Syntax (EX Series and QFX5100 Switch) on page 198
Syntax (QFX10002 Switch) on page 198

Syntax (EX Series and QFX5100 Switch)

```
layer2 {
  no-destination-mac-address;
  no-ether-type;
  no-source-mac-address;
  vlan-id;
}
```

Syntax (QFX10002 Switch)

```
layer2 {
  no-incoming-port;
  no-destination-mac-address;
  no-ether-type;
  no-source-mac-address;
  source-mac-address;
  vlan-id;
  no-vlan-id;
  inner-vlan-id;
}
```

Hierarchy Level

[edit forwarding-options enhanced-hash-key]

Release Information

Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.

Description

Select the fields in the Layer 2 header that are used by the hashing algorithm to make hashing decisions.

When traffic enters a link aggregation group (LAG) bundle, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links. The hashing algorithm always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.

When traffic is exiting a device that has enabled ECMP, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to forward traffic to the next hop device.

The hashing algorithm only inspects the fields in the Layer 2 header when the hash mode is set to Layer 2 header. You can set the hash mode to Layer 2 header using the **set forwarding-options enhanced-hash-key hash-mode layer2-header** statement.

Default

The hash mode of the hashing algorithm is set to Layer 2 payload, by default. When the hash mode is set to Layer 2 payload, the hashing algorithm does not use fields in the Layer 2 header to make hashing decisions.

The following fields are used by the hashing algorithm when the hash mode of the hashing algorithm is set to Layer 2 header, by default:

- Destination MAC address
- Ethertype
- Source MAC address

Options

no-destination-mac-address—Exclude the destination MAC address field from the hashing algorithm.

no-ether-type—Exclude the Ethertype field from the hashing algorithm.

no-source-mac-address—Exclude the source MAC address field from the hashing algorithm.

vlan-id—Include the VLAN ID field in the hashing algorithm.

Required Privilege Level

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

Related Documentation

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 82
- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 62
- **enhanced-hash-key on page 165**
- **hash-mode on page 180**

link-mode

Syntax link-mode *mode* (automatic | full-duplex | half-duplex);

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 can operate in either full-duplex or half-duplex mode. The router's or switch's management Ethernet interface, **fxp0** or **em0**, and the built-in Fast Ethernet interfaces on the FIC (M7i router) 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 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)*

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25

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.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.
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*]** hierarchy level or at the **[edit interfaces xe-*fpc/pic/port* ether-options 802.3ad ae*x*]** hierarchy level.

  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 remaining statements are explained separately. See CLI Explorer.

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 78
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73

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.
mixed option added in Junos OS Release 15.1F3 and 16.1R2 for PTX5000 routers and 15.1F6 and 16.1R2 for PTX3000 routers.

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, MX Series, PTX Series routers, and QFX5100, QFX10002, QFX10008, and QFX10016 switches 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.

mixed—Enables bundling of different Ethernet rate links in the same 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

- *Aggregated Ethernet Interfaces Overview*
- Configuring Aggregated Ethernet Link Speed on page 79
- *Configuring Mixed Rates and Mixed Modes on Aggregated Ethernet Bundles*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

liveness-detection

Syntax

```
liveness-detection {
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
```

Hierarchy Level [edit protocols iccp peer]

Release Information

Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Enable Bidirectional Forwarding Detection (BFD). BFD enables rapid detection of communication failures between peers.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

local-bias

Syntax local-bias;

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.

Description Enable local link bias for all links in the aggregated Ethernet interface.

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.

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 111
- Understanding Local Link Bias on page 109

local-ip-addr (ICCP)

Syntax local-ip-addr *local-ip-address*;

Hierarchy Level [edit protocols iccp],
[edit protocols iccp peer *peer-IP-address*]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the local IP address of the interchassis link (ICL) interface that Inter-Chassis Control Protocol (ICCP) uses to communicate to the peers that host a multichassis link aggregation group (MC-LAG).

Options ***local-ip-address***—Default local IP address to be used by all peers.

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

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 51

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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- *Understanding Interface Ranges on EX Series Switches*
- Understanding Interface Ranges on EX Series Switches on page 89
- EX Series Switches Interfaces Overview on page 19
- *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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- *Understanding Interface Ranges on EX Series Switches*
- Understanding Interface Ranges on EX Series Switches on page 89
- EX Series Switches Interfaces Overview on page 19
- *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*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 30
- *Configuring VLANs for EX Series Switches (CLI Procedure)*
- *Configuring VLANs for EX Series Switches (CLI Procedure)*

minimum-interval (Liveness Detection)

Syntax minimum-interval *milliseconds*;

Hierarchy Level [edit protocols iccp peer liveness-detection]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure simultaneously the minimum interval at which the peer transmits liveness detection requests and the minimum interval at which the peer expects to receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately by using the **transmit-interval minimal-interval** and **minimum-receive-interval** statements, respectively.

Options *milliseconds*—Specify the minimum interval value for Bidirectional Forwarding Detection (BFD).

Range: 1 through 255,000

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

minimum-receive-interval (Liveness Detection)

Syntax minimum-receive-interval *milliseconds*;

Hierarchy Level [edit protocols iccp peer liveness-detection]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the minimum interval at which the peer must receive a reply from a peer with which it has established a Bidirectional Forwarding Detection (BFD) session.

Options *milliseconds*—Specify the minimum interval value.

Range: 1 through 255,000

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

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 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.
- Starting in Release 14.2, MTU for IRB interfaces is calculated by removing the Ethernet header overhead (**6(DMAC)+6(SMAC)+2(EtherType)**), and the MTU is a minimum of the two values:
  - Configured MTU
  - Associated bridge domain's physical or logical interface MTU
    - For Layer 2 logical interfaces configured with **flexible-vlan-tagging**, IRB MTU is calculated by including 8 bytes overhead (**SVLAN+CVLAN**).
    - For Layer 2 logical interfaces configured with **vlan-tagging**, IRB MTU is calculated by including single VLAN 4 bytes overhead.

NOTE: Changing the Layer 2 logical interface option from **vlan-tagging** to **flexible-vlan-tagging** or vice versa adjusts the logical interface MTU by 4 bytes with the existing MTU size. As a result, the Layer 2 logical interface is deleted and re-added, and the IRB MTU is re-computed appropriately.

For more information about configuring MTU for specific interfaces and router or switch combinations, see *Configuring the Media MTU*.

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), 256 through 9500 bytes (Junos OS 16.1R1 for MX Series routers)

NOTE: Starting in Junos OS Release 16.1R1, the MTU size for a media or protocol is increased from 9192 to 9500 for Ethernet interfaces on the following MX Series MPCs:

- MPC1
- MPC2
- MPC2E
- MPC3E
- MPC4E
- MPC5E
- MPC6E

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 the Media MTU*
- *Configuring the MTU for Layer 2 Interfaces*
- *Setting the Protocol MTU*

native-vlan-id

Syntax native-vlan-id *number*;

Hierarchy Level [edit interfaces ge-*fpc/pic/port*],
[edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Statement introduced in Junos OS Release 13.2X51-D20 for the QFX Series.

Description Configure mixed tagging support for untagged packets on a port for the following:

- M Series routers with Gigabit Ethernet IQ PICs with SFP and Gigabit Ethernet IQ2 PICs with SFP configured for 802.1Q flexible VLAN tagging
- MX Series routers with Gigabit Ethernet DPCs and MICs, Tri-Rate Ethernet DPCs and MICs, and 10-Gigabit Ethernet DPCs and MICs and MPCs configured for 802.1Q flexible VLAN tagging
- T4000 routers with 100-Gigabit Ethernet Type 5 PIC with CFP
- EX Series switches with Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces

When the **native-vlan-id** statement is included with the *flexible-vlan-tagging* statement, untagged packets are accepted on the same mixed VLAN-tagged port.

NOTE: **The logical interface on which untagged packets are received must be configured with the same VLAN ID as the native VLAN ID configured on the physical interface, otherwise the untagged packets are dropped.**

To configure the logical interface, include the **vlan-id** statement (matching the **native-vlan-id** statement on the physical interface) at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

When the **native-vlan-id** statement is included with the interface-mode statement, untagged packets are accepted and forwarded within the bridge domain or VLAN that is configured with the matching VLAN ID.

Starting in Junos OS Release 17.1R1, you can send untagged traffic without a native VLAN ID to the remote end of the network. To do this, remove the native VLAN ID from the untagged traffic configuration by setting the **no-native-vlan-insert** statement. If you do not configure this statement, the native VLAN ID is added to the untagged traffic.

Default By default, the untagged packets are dropped. That is, if you do not configure the **native-vlan-id** option, the untagged packets are dropped.

Options *number*—VLAN ID number.
Range: (ACX Series routers and EX Series switches) 0 through 4094.

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

Related Documentation
- *Configuring Mixed Tagging Support for Untagged Packets*
- *Configuring Access Mode on a Logical Interface*
- *Configuring the Native VLAN Identifier (CLI Procedure)*
- *Understanding Bridging and VLANs on EX Series Switches*
- *flexible-vlan-tagging*
- *Understanding Q-in-Q Tunneling on EX Series Switches*
- *no-native-vlan-insert*
- *Sending Untagged Traffic Without VLAN ID to Remote End*

no-gratuitous-arp-request

Syntax no-gratuitous-arp-request;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.

Description For Ethernet interfaces and pseudowire logical interfaces, do not respond to gratuitous ARP requests.

Default Gratuitous ARP responses are enabled 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 52
- gratuitous-arp-reply on page 179

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 56
- *Junos OS Administration Library*

peer (ICCP)

Syntax

```
peer ip-address {
  authentication-key string;
  backup-liveness-detection {
    backup-peer-ip ip-address;
  }
  liveness-detection {
    detection-time {
      threshold milliseconds;
    }
    minimum-intervalmilliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  local-ip-addr ipv4-address;
  session-establishment-hold-time seconds;
}
```

Hierarchy Level [edit protocols iccp]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the peers that host a multichassis link aggregation group (MC-LAG). You must configure Inter-Chassis Control Protocol (ICCP) for both peers that host the MC-LAG.

NOTE: Backup liveness detection is not supported on MX Series routers.

The remaining statements are explained separately. See CLI Explorer.

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

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.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.

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 72
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*

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 46

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 46

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 57
- *Configuring Proxy ARP (CLI Procedure)*
- *Example: Configuring Proxy ARP on an EX Series Switch*
- Configuring Gratuitous ARP on page 52

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*
- Configuring Unicast RPF (CLI Procedure) on page 117
- Disabling Unicast RPF (CLI Procedure) on page 119
- Understanding Unicast RPF on page 113

session-establishment-hold-time

Syntax session-establishment-hold-time *seconds*;

Hierarchy Level [edit protocols iccp],
[edit protocols iccp peer]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Specify the time during which an Inter-Chassis Control Protocol (ICCP) connection must be established between peers.

Options ***seconds***—Time (in seconds) within which a successful ICCP connection must be established.

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

speed (Ethernet)

Syntax speed (10m | 100m | 1g | auto | auto-10m-100m);

Hierarchy Level [edit interfaces *interface-name*],
[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 12.2 for ACX Series Universal Access Routers.
Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.

Description Configure the interface speed. This statement applies to the management Ethernet interface (**fxp0** or **em0**), Fast Ethernet 12-port and 48-port PICs, the built-in Fast Ethernet port on the FIC (M7i router), Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers,and Gigabit Ethernet interfaces on EX Series switches.

When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled. When you configure 100BASE-FX SFP, you must set the port speed at 100 Mbps.

NOTE: On MX Series routers with Tri-rate Enhanced DPC (DPCE-R-40GE-TX), when you configure the interface speed using the **auto-10m-100m** option, the speed is negotiated to the highest value possible (100 Mbps), if the same value is configured on both sides of the link. However, when you view the interface speed of the DPC, using the **show interfaces** command, the value of the speed is not accurately displayed. For instance, if you configure the speed of the Tri-rate enhanced DPC, as 100Mbps on both sides of the link, the interface speed of the DPC is negotiated to 100 Mbps. However, the interface speed of the DPC displays 1 bps. This is an issue with the **show interfaces** command only. The actual interface speed is 100 Mbps.

Options You can specify the speed as either **10m** (10 Mbps), **100m** (100 Mbps), and on MX Series routers, **1g** (1 Gbps). You can also specify the **auto** option on MX Series routers.

For Gigabit Ethernet interfaces on EX Series switches, you can specify one of the following options:

- **10m**—10 Mbps
- **100m**—100 Mbps
- **1g**—1 Gbps
- **auto**—Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link.

- **auto-10m-100m**—Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end 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 the Interface Speed*
- *Configuring the Interface Speed on Ethernet Interfaces*
- *Configuring Gigabit Ethernet Autonegotiation*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25

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 125

transmit-interval (Liveness Detection)

Syntax

```
transmit-interval {
  minimum-interval milliseconds;
  threshold milliseconds;
}
```

Hierarchy Level [edit protocols iccp peer liveness-detection]

Release Information Statement introduced in Junos OS Release 10.0 for MX Series routers.
Statement introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.

Description Configure the Bidirectional Forwarding Detection (BFD) transmit interval. The negotiated transmit interval for a peer is the interval between the sending of BFD liveness detection requests to peers. The receive interval for a peer is the minimum interval between receiving packets sent from its peer; the receive interval is not negotiated between peers. To determine the transmit interval, each peer compares its configured minimum transmit interval with its peer's minimum receive interval. The larger of the two numbers is accepted as the transmit interval for that peer.

The remaining statements are explained separately. See CLI Explorer.

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

traps

Syntax (traps | no-traps);

Hierarchy Level [edit dynamic-profiles *profile-name* interfaces *interface-name*],
[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.
Support at the **[edit dynamic-profiles *profile-name* interfaces*interface-name*]**hierarchy level introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.

Description Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.

(Enhanced subscriber management for MX Series routers) To enable SNMP notifications, you must first configure the **interface-mib** statement at the **[edit dynamic-profiles *profile-name* interfaces *interface-name*]** hierarchy level. If **interface-mib** is not configured, the **traps** statement has no effect.

BEST PRACTICE: To achieve maximum performance when enhanced subscriber management is enabled, we recommend that you *not* enable SNMP notifications on all dynamic subscriber interfaces.

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*
- Enabling or Disabling SNMP Notifications on Logical Interfaces on page 57

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. See CLI Explorer.

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)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- EX Series Switches Interfaces Overview on page 19
- *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. See CLI Explorer.

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 238
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 25
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 30
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 106
- *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.

NOTE: On EX Series switches except for EX4300 and EX9200 switches, the **vlan-tagging** and **family ethernet-switching** statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default **family** setting.

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

Related Documentation

- 802.1Q VLANs Overview on page 105
- **vlan-id on page 237**
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 106
- Configuring Tagged Aggregated Ethernet Interfaces on page 85
- *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*

CHAPTER 12

Operational Commands

- monitor interface
- request diagnostics tdr
- show diagnostics tdr
- show forwarding-options enhanced-hash-key
- show interfaces diagnostics optics
- show interfaces ge-
- show interfaces irb
- show interfaces mc-ae
- show interfaces me0
- show interfaces queue
- show interfaces xe-
- show lacp interfaces
- test interface restart-auto-negotiation

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 switch.

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 21 on page 240. The keys are not case-sensitive.

Table 21: Output Control Keys for the monitor interface interface-name 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.

Table 21: Output Control Keys for the monitor interface interface-name Command *(continued)*

Key	Action
i	Displays information about a different interface. The command prompts you for the name of a specific interface.
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 22 on page 241. The keys are not case-sensitive.

Table 22: Output Control Keys for the monitor interface traffic Command

Key	Action
b	Displays the statistics in units of bytes and bytes 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 243
monitor interface (OTN Interface) on page 244
monitor interface (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface) on page 245
monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface) on page 246
monitor interface (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface) on page 247
monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface) on page 247

monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface) on page 248
monitor interface (Logical) on page 249
monitor interface (QFX3500 Switch) on page 249
monitor interface traffic on page 250
monitor interface traffic (QFX3500 Switch) on page 250
monitor interface traffic detail (QFX3500 Switch) on page 251

Output Fields Table 23 on page 242 describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 23: 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

Table 23: monitor interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
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
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 (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-0/0/3
Interface: xe-0/0/3, 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:                5                                  [0]
  Output errors:                      0                                  [0]
  Output drops:                       0                                  [0]
  Aged packets:                       0                                  [0]
Active alarms : None
Active defects: None
PCS statistics:                     Seconds
  Bit Errors                          0                                  [0]
  Errored blocks                      4                                  [0]
Input MAC/Filter statistics:
  Unicast packets                     0                                  [0]
  Broadcast packets                   0                                  [0]
  Multicast packets                   0                                  [0]
```

```
  Oversized frames                      0                                [0]
  Packet reject count                   0                                [0]
  DA rejects                            0                                [0]
  SA rejects                            0                                [0]
Output MAC/Filter Statistics:
  Unicast packets                       0                                [0]
  Broadcast packets                     0                                [0]
  Multicast packets                     0                                [0]
  Packet pad count                      0                                [0]
  Packet error count                    0                                [0]



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

monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface)

```
user@host> monitor interface et-2/1/0
Interface: et-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
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:                263                                [0]
  Output errors:                        0                                [0]
  Output drops:                         0                                [0]
  Aged packets:                         0                                [0]
OTN Link 0
  OTN Alarms:
  OTN Defects:
  OTN OC - Seconds
    LOS                               129                                [0]
    LOF                                 2                                [0]
  OTN OTU - FEC Statistics
    Corr err ratio                  <8E-5
    Corr bytes              169828399453                                 [0]
    Uncorr words             28939961456                                 [0]
  OTN OTU - Counters                                                     [0]
    BIP                                 0
    BBE                                 0                                [0]
    ES                                 24                                [0]
    SES                                 0                                [0]
    UAS                              1255                                [0]
  OTN ODU - Counters                                                     [0]
    BIP                                 0
    BBE                                 0                                [0]
    ES                                 24                                [0]
    SES                                 0                                [0]
    UAS                              1256                                [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 (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-6/1/0
Interface: xe-6/1/0, 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:                1                                  [0]
  Output errors:                      0                                  [0]
  Output drops:                       0                                  [0]
  Aged packets:                       0                                  [0]
Active alarms : None
Active defects: None
PCS statistics:                     Seconds
  Bit Errors                          0                                  [0]
  Errored blocks                      1                                  [0]
Input MAC/Filter statistics:
  Unicast packets                     0                                  [0]
  Broadcast packets                   0                                  [0]
  Multicast packets                   0                                  [0]
  Oversized frames                    0                                  [0]
  Packet reject count                 0                                  [0]
  DA rejects                          0                                  [0]
  SA rejects                          0                                  [0]
Output MAC/Filter Statistics:
  Unicast packets                     0                                  [0]
  Broadcast packets                   0                                  [0]
  Multicast packets                   0                                  [0]
  Packet pad count                    0                                  [0]
  Packet error count                  0                                  [0]


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

monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface et-9/0/0
Interface: et-9/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
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:             1                                  [0]
  Output errors:                   0                                  [0]
  Output drops:                    0                                  [0]
  Aged packets:                    0                                  [0]



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

monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```
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 254
- Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) on page 138

List of Sample Output request diagnostics tdr start interface ge-0/0/19 on page 253

Output Fields Table 24 on page 253 lists the output fields for the **request diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 24: 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 252
- Diagnosing a Faulty Twisted-Pair Cable (CLI Procedure) on page 138

List of Sample Output

show diagnostics tdr interface ge-0/0/19 (Normal Cable) on page 256
show diagnostics tdr interface ge-2/0/2 (Faulty Cable) on page 257
show diagnostics tdr (All Supported Interfaces) on page 257

Output Fields

Table 25 on page 255 lists the output fields for the **show diagnostics tdr** command. Output fields are listed in the approximate order in which they appear.

Table 25: 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 25: 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 forwarding-options enhanced-hash-key

Syntax **show forwarding-options enhanced-hash-key**

Release Information Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.
Fabric Load Balancing Options output fields introduced in Junos OS Release 14.1X53-D10.

Description Display information about which packet fields are used by the hashing algorithm to make hashing decisions.

You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the **forwarding-options** enhanced-hash-key statement.

Required Privilege Level view

Related Documentation

- Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 82
- Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 62
- enhanced-hash-key on page 165

List of Sample Output

show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode) on page 261
show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode) on page 261
show forwarding-options enhanced-hash-key (Fabric Load Balancing Options) on page 262
show forwarding-options enhanced-hash-key (QFX10002 and QFX 10008 Switches) on page 262

Output Fields Table 26 on page 259 lists the output fields for the **show forwarding-options enhanced-hash-key** command. Output fields are listed in the approximate order in which they first appear. Output fields vary by platform.

Table 26: show forwarding-options enhanced-hash-key Output Fields

Field Name	Field Description
Hash-Mode	Current hash mode: Layer 2 header or Layer 2 payload.
Protocol	Indicates whether the Protocol field is or is not used by the hashing algorithm: Yes or No.
Destination L4 Port	Indicates whether the Destination L4 Port field is or is not used by the hashing algorithm: Yes or No.

Table 26: show forwarding-options enhanced-hash-key Output Fields *(continued)*

Field Name	Field Description
Source L4 Port	Indicates whether the Source L4 Port field is or is not used by the hashing algorithm: Yes or No.
Destination IPv4 Addr	Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
Source IPv4 Addr	Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No.
Vlan id	Indicates whether the Vlan ID field is or is not used by the hashing algorithm: Yes or No.
Inner-Vlan ID	indicates whether the inner Vlan field is or is not used by the hashing algorithnm: Yes or No.
Next Hdr	Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No.
Destination IPv6 Addr	Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
Source IPv6 Addr	Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No.
Ether Type	Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No.
Destination MAC Address	Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No.
Source MAC Address	Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No.
Load Balancing Method	Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet. The load balancing method is flowlet by default, and can be configured using the **fabric-load-balance** statement.
Fabric Link Scale	Indicates the fabric link scale, in mbps.
Inactivity Interval	Indicates the fabric load balance inactivity interval, in microseconds (us). The inactivity interval is 16 microseconds by default, and can be configured using the **inactivity-interval** statement.
Hash Region Size/Trunk	Indicates the hash region size, in buckets per fabric trunk.

Table 26: show forwarding-options enhanced-hash-key Output Fields *(continued)*

Field Name	Field Description
Seed	A hash seed value, between 0 and 4294967295. If a hash-seed value is not configured it is automatically assigned on the QFX10002 and QFX10008 switches. A hash-seed prevents traffic polarization to same links on the next hop QFX switch when two are connected with LAG/ECMP.
Key	Indicates whether the GRE key field is or is not used by the hashing algorithm: Yes or No.
Protocol	Indicates if a Generic Router Encapsulation (GRE) endpoint over routes was dynamically learned by a routing protocol such as RIP or OSPF.
MPLS Enabled	Indicates if MPLS is enabled under L2 switching.
VXLAN VNID	A 24-bit virtual network identifier (VNID) that uniquely identifies the Virtual Extensible Local Area Networks (VXLAN) segment.

Sample Output

show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0


Current Hash Settings
---------------------
Hash-Mode                              :layer2-payload

  inet Hash settings-
  -------------------
    inet packet fields
       Protocol                        : Yes
       Destination L4 Port             : Yes
       Source L4 Port                  : Yes
       Destination IPv4 Addr           : Yes
       Source IPv4 Addr                : Yes
       Vlan id                         : No


  inet6 Hash settings-
  --------------------
    inet6 packet fields
       Next Hdr                        : Yes
       Destination L4 Port             : Yes
       Source L4 Port                  : Yes
       Destination IPv6 Addr           : Yes
       Source IPv6 Addr                : Yes
       Vlan id                         : No
```

show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode)

```
user@switch> show forwarding-options enhanced-hash-key
```

```
Slot 0


Current Hash Settings
---------------------

Hash-Mode                              : layer2-header

  layer2 Hash settings-
  ---------------------
     layer2 packet fields
       Ether Type                      : Yes
       Destination MAC Address         : Yes
       Source MAC Address              : Yes
       VLAN ID                         : No
```

show forwarding-options enhanced-hash-key (Fabric Load Balancing Options)

```
user@switch> show forwarding-options enhanced-hash-key
<some output removed for brevity>

  Fabric Load Balancing Options
  ------------------------------------------------------------
      Load Balancing Method  : Flowlet
      Fabric Link Scale      : 40960 (mbps)
      Inactivity Interval    : 16 (us)
      Hash Region Size/Trunk : 1024 (buckets)
```

show forwarding-options enhanced-hash-key (QFX10002 and QFX 10008 Switches)

```
user@switch> show forwarding-options enhanced-hash-key
Slot 0


Seed value for Hash function           0: 3626023417
Seed value for Hash function           1: 3626023417
Seed value for Hash function           2: 3626023417
Seed value for Hash function           3: 3626023417

    Inet settings:
    --------------
        IPV4 dest address:    Yes
        IPV4 source address:  Yes
        L4 Dest Port:         Yes
        L4 Source Port:       Yes
    Inet6 settings:
    --------------
        IPV6 dest address:    Yes
        IPV6 source address:  Yes
        L4 Dest Port:         Yes
        L4 Source Port:       Yes
    L2 settings:
    ------------
        Dest Mac address:     No
        Source Mac address:   No
        Vlan Id:              Yes
        Inner-vlan Id:        No
        Incoming port:        Yes
    GRE settings:
    -------------
```

```
    Key:                No
    Protocol:           No
MPLS settings:
--------------
    MPLS Enabled:       Yes
VXLAN settings:
---------------
    VXLAN VNID:         No
```

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 123
- *Monitoring Interface Status and Traffic*
- *Installing a Transceiver*
- *Installing a Transceiver in a QFX Series Device*
- *Removing a Transceiver*
- *Removing a Transceiver from a QFX Series Device*
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

Output Fields Table 27 on page 265 lists the output fields for the **show interfaces diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

Table 27: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
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 27: 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 27: 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 27: 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 27: 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 27: 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 27: 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 123
- *Troubleshooting Network Interfaces on EX3200 Switches*
- *Troubleshooting Network Interfaces on EX4200 Switches*
- Troubleshooting an Aggregated Ethernet Interface on page 134
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

show interfaces ge-0/0/0 on page 285
show interfaces ge-0/0/0 brief on page 285
show interfaces ge-0/0/0 brief (with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces) on page 286
show interfaces ge-0/0/0 detail on page 286
show interfaces ge-0/0/4 extensive on page 287

Output Fields

Table 28 on page 279 lists the output fields for the **show interfaces ge-** command. Output fields are listed in the approximate order in which they appear.

Table 28: 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 28: 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 28: 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 28: 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 28: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	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).
		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

Logical Interface		

Table 28: 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 28: 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 irb

Syntax

show interfaces irb
<brief | detail | extensive | terse>
<descriptions>
<media>
<routing-instance *instance-name*>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 12.3R2.
Command introduced in Junos OS Release 12.3R2 for EX Series switches.
Command introduced in Junos OS Release 13.2 for the QFX Series

Description

Display integrated routing and bridging interfaces information.

Options

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 *instance-name*—(Optional) Display information for the interface with the specified SNMP index.

snmp-index *snmp-index*—(Optional) Display information for the interface with the specified SNMP index.

statistics—(Optional) Display static interface statistics.

Additional Information

Integrated routing and bridging (IRB) provides simultaneous support for Layer 2 bridging and Layer 3 IP routing on the same interface. IRB enables you to route local packets to another routed interface or to another VLAN that has a Layer 3 protocol configured.

Required Privilege Level

view

List of Sample Output

show interfaces irb extensive on page 294
show interfaces irb snmp-index on page 295

Output Fields

Table 29 on page 290 lists the output fields for the **show interfaces irb** command. Output fields are listed in the approximate order in which they appear.

Table 29: show interfaces irb Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 29: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
Enabled	State of the physical interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Proto	Protocol configured on the interface.	**terse**
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Type	Physical interface type.	**detail extensive** none
Link-level type	Encapsulation being used on the physical interface.	**detail extensive brief** none
MTU	MTU size on the physical interface.	**detail extensive brief** none
Clocking	Reference clock source: **Internal** or **External**. Always unspecified on IRB interfaces.	**detail extensive brief**
Speed	Speed at which the interface is running. Always unspecified on IRB interfaces.	**detail extensive brief**
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	**detail extensive brief** none
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	**detail extensive brief** none
Link type	Physical interface link type: **full duplex** or **half duplex**.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	**detail extensive** none
Physical Info	Physical interface information.	All levels
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	Backup address of the link.	**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 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** none

Table 29: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
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.	**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**
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. • **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 not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **Resource errors**—Sum of transmit drops.	**detail 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, increasing only when the cable is unplugged, the far-end system is powered down and 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 DPC 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**

Table 29: show interfaces irb 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 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" section under *Common Output Fields Description*.	**detail extensive**
Encapsulation	Encapsulation on the logical interface.	**detail extensive**
Bandwidth	Speed at which the interface is running.	**detail extensive**
Routing Instance	Routing instance IRB is configured under.	**detail extensive**
Bridging Domain	Bridging domain IRB is participating in.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical 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**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical 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**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router.	**detail extensive**
Protocol	Protocol family configured on the local interface. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive**
MTU	Maximum transmission unit size on the logical interface.	**detail extensive**

Table 29: show interfaces irb Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
Addresses, Flags	Information about address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive**
Policer	The policer that is to be evaluated when packets are received or transmitted on the interface.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	**detail extensive**

Sample Output

show interfaces irb extensive

```
user@host> show interfaces irb extensive
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23, Generation: 130
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
Speed: Unspecified
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  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 irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
    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, MTU: 1500, Generation: 154, Route table: 0
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/8, Local: 10.51.1.2, Broadcast: 10.51.1.255,
        Generation: 155
    Protocol multiservice, MTU: 1500, Generation: 155, Route table: 0
      Flags: Is-Primary
      Policer: Input: __default_arp_policer
```

show interfaces irb snmp-index

```
user@host> show interfaces irb snmp-index 25
Physical interface: irb, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 25
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Current address: 02:00:00:00:00:30, Hardware address: 02:00:00:00:00:30
  Last flapped   : Never
    Input packets : 0
    Output packets: 0

  Logical interface irb.0 (Index 68) (SNMP ifIndex 70)
    Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
    Bandwidth: 1000mbps
    Routing Instance: customer_0 Bridging Domain: bd0
    Input packets : 0
    Output packets: 0
    Protocol inet, MTU: 1500
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.1/8, Local: 10.51.1.2, Broadcast: 10.51.1.255
```

```
    Protocol multiservice, MTU: 1500
      Flags: Is-Primary
```

show interfaces mc-ae

Syntax **show interfaces mc-ae id** *identifier* **unit** *number*

Release Information Command introduced in Junos OS Release 9.6 for the MX Series.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Statement introduced in Junos OS Release 12.3R2 for EX Series switches.
Configuration Consistency Check output field added in Junos OS Release 15.1X53-D60 for the QFX Series.

Description On peers with multichassis aggregated Ethernet (**mc-ae***X*) interfaces, use this command to display information about the multichassis aggregated Ethernet interfaces.

Options **id** *identifier*—(Optional) Specify the name of the multichassis aggregated Ethernet interface.

unit *number*—(Optional) Specify the logical interface by unit number.

Required Privilege Level view

List of Sample Output
show interfaces mc-ae (EX Series) on page 298
show interfaces mc-ae (MX Series) on page 298
show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series) on page 298

Output Fields Table 30 on page 297 lists the output fields for the **show interfaces mc-ae** command. Output fields are listed in the approximate order in which they appear.

Table 30: show interfaces mc-ae Output Fields

Output Field Name	Field Description
Current State Machine's State	Specifies the state of the MC-LAG initialization state machine.
Configuration Consistency Check	Specifies the status of the MC-LAG configuration consistency check feature. The status is either **Passed** or **Failed**. If the status is **Failed**, the system will display the name of the parameter that failed consistency check. If there are multiple inconsistencies, only the first inconsistency is shown. If the enforcement level for the MC-LAG parameter was mandatory, and you did not configure that parameter correctly, the command will show that the MC-LAG interface is down.
Member Link	Specifies the identifiers of the configured multichassis link aggregated interface members.
Local Status	Specifies the status of the local link: **active** or **standby**.
Peer Status	Specifies the status of the peer link: **active** or **standby**.

Table 30: show interfaces mc-ae Output Fields *(continued)*

Output Field Name	Field Description
Peer State	Specifies the status of the local and peer links in an **active/active** MC-LAG configuration.
Logical Interface	Specifies the identifier and unit of the AE interface.
Topology Type	Specifies the bridge configured on the AE.
Local State	Specifies if the local device is up or down.
Peer State	Specifies if the peer device is up or down.
Peer Ip/MCP/State	Specifies the multichassis protection (MCP) link or the interchassis link-protection link (ICL-PL) for all of the multichassis aggregated Ethernet interfaces that are part of the peer.

Sample Output

show interfaces mc-ae (EX Series)

```
user@switch> show interfaces mc-ae ae1 512
 Member Link                  : ae1
  Current State Machine's State: mcae active state
  Configuration Consistency Check    : Failed (redundancy group id mismatch)
  Local Status                 : active
  Local State                  : up
  Peer Status                  : standby
  Peer State                   : up
      Logical Interface        : ae1.0
      Topology Type            : bridge
      Local State              : up
      Peer State               : up
      Peer Ip/MCP/State        : 10.1.1.1 ae0.0 up
```

show interfaces mc-ae (MX Series)

```
user@host> show interfaces mc-ae ae0 unit 512
  Member Links   : ae0
  Local Status   : active
  Peer Status    : active
  Logical Interface        : ae0.512
  Core Facing Interface : Label Ethernet Interface
  ICL-PL         : Label Ethernet Interface
```

show interfaces mc-ae (Active/Active Bridging and VRRP over IRB on MX Series)

```
user@host# show interfaces mc-ae ge-0/0/0.0
 Member Link                  : ae0
 Current State Machine's State: active
 Local Status                 : active
 Local State                  : up
 Peer Status                  : active
```

```
Peer State                  : up
    Logical Interface       : ae0.0
    Topology Type           : bridge
    Local State             : up
    Peer State              : up
    Peer Ip/ICL-PL/State    : 192.168.100.10 ge-0/0/0.0 up
```

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 304
show interfaces me0 brief on page 304
show interfaces me0 detail on page 304
show interfaces me0 extensive on page 305

Output Fields

Table 31 on page 300 lists the output fields for the **show interfaces me0** command. Output fields are listed in the approximate order in which they appear.

Table 31: 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

Table 31: show interfaces me0 Output Fields *(continued)*

Field Name	Field Description	Level of Output
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**
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*** **(*weeks*w:*days*d*hour: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*** **(*weeks*w:*days*d*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3w:3d 22:58 ago).**	**detail extensive**

Table 31: 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 31: 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 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 123
- *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 309
show interfaces queue xe-6/0/39 (Line Card with Oversubscribed Ports in an EX8200 Switch) on page 310

Output Fields

Table 32 on page 308 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 32: show interfaces queue Output Fields

Field Name	Field Description
Physical Interface and Forwarding Class Information	
Physical interface	Name of the physical interface.
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 32: 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 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 123
- *Troubleshooting Network Interfaces on EX3200 Switches*
- *Troubleshooting Network Interfaces on EX4200 Switches*

- Troubleshooting an Aggregated Ethernet Interface on page 134
- *Junos OS Ethernet Interfaces Configuration Guide*

List of Sample Output

show interfaces xe-4/1/0 on page 322
show interfaces xe-0/1/0 brief on page 323
show interfaces xe-4/1/0 detail on page 323
show interfaces xe-6/0/39 extensive on page 324

Output Fields

Table 33 on page 314 lists the output fields for the **show interfaces xe-** command. Output fields are listed in the approximate order in which they appear.

Table 33: 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

Table 33: show interfaces xe- 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	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

Table 33: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flow control	Flow control status: **Enabled** or **Disabled**.	**brief** **detail** **extensive** none
Device flags	Information about the physical device.	**brief** **detail** **extensive** none
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**

Table 33: show interfaces xe- 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. • **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**
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 33: 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 33: 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 33: 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 33: 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 33: 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 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.
Command introduced in Junos OS Release 14.2R3

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*
- *Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch*
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- Configuring Aggregated Ethernet Links (CLI Procedure) on page 68
- *Configuring Link Aggregation*
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 72
- *Configuring Aggregated Ethernet LACP*
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) on page 73
- Understanding Aggregated Ethernet Interfaces and LACP on page 59
- *Understanding Aggregated Ethernet Interfaces and LACP*
- *Junos OS Interfaces Fundamentals Configuration Guide*

List of Sample Output

show lacp interfaces (EX Series Switches) on page 329

show lacp interfaces (QFX Series) on page 330

Output Fields Table 34 on page 328 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 34: 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 34: 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
```

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 332

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
```