

Interfaces Feature Guide



Modified: 2019-05-14



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

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

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

Interfaces Feature Guide

Copyright © 2019 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

| | | |
|------------------|---|----------|
| | About the Documentation | xiii |
| | Documentation and Release Notes | xiii |
| | Using the Examples in This Manual | xiii |
| | Merging a Full Example | xiv |
| | Merging a Snippet | xiv |
| | Documentation Conventions | xv |
| | Documentation Feedback | xvii |
| | Requesting Technical Support | xvii |
| | Self-Help Online Tools and Resources | xviii |
| | Creating a Service Request with JTAC | xviii |
| Part 1 | Ethernet OAM Link Fault Management | |
| Chapter 1 | Understanding Ethernet OAM Link Fault Management | 3 |
| | Understanding Ethernet OAM Link Fault Management | 3 |
| | Configuring Ethernet OAM Link Fault Management | 4 |
| Part 2 | Interfaces | |
| Chapter 2 | Understanding Interfaces | 9 |
| | Interfaces Overview for Switches | 9 |
| | Network Interfaces for EX Series | 9 |
| | Special Interfaces for EX Series | 10 |
| | Network Interfaces for EX4600, NFX Series, QFX Series, QFabric System | 12 |
| | Special Interfaces for EX4600, NFX Series, QFX Series, QFabric System | 13 |
| | Network Interfaces for OCX Series | 13 |
| | Special Interfaces for OCX Series | 14 |
| | Understanding Interface Naming Conventions | 14 |
| | Physical Part of an Interface Name for EX Series | 15 |
| | Logical Part of an Interface Name for EX Series | 16 |
| | Wildcard Characters in Interface Names for EX Series | 17 |
| | Physical Part of an Interface Name for QFX series, NFX Series, EX4600, QFabric System | 17 |
| | Logical Part of an Interface Name on a Switch Running QFabric Software Package for QFX series, NFX Series, EX4600, QFabric System | 25 |
| | Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software for QFX series, NFX Series, EX4600, QFabric System | 25 |
| | Wildcard Characters in Interface Names for QFX series, NFX Series, EX4600, QFabric System | 26 |

| | | |
|------------------|--|-----------|
| | Physical Part of an Interface Name for OCX1100 | 26 |
| | Wildcard Characters in Interface Names for OCX1100 | 26 |
| | Understanding Interface Ranges for Switches | 27 |
| | Understanding Management Interfaces | 29 |
| | Understanding Port Ranges and System Modes | 30 |
| | Port Ranges for Different Media Types | 31 |
| | Supported System Modes | 54 |
| | Configuring the Interface Address | 57 |
| | Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches | 59 |
| | Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 60 |
| | Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches | 61 |
| | Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches | 61 |
| | Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches | 63 |
| | Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches | 63 |
| | Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 64 |
| | Configuring Ethernet Loopback Capability | 65 |
| | Monitoring Interface Status and Traffic | 65 |
| | Troubleshooting Network Interfaces | 65 |
| | The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down | 65 |
| Part 3 | Layer 3 Logical Interfaces | |
| Chapter 3 | Understanding Layer 3 Logical Interfaces | 69 |
| | Understanding Layer 3 Logical Interfaces | 69 |
| | Configuring a Layer 3 Logical Interface | 70 |
| | Verifying That Layer 3 Logical Interfaces Are Working | 70 |
| Part 4 | Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP) | |
| Chapter 4 | Understanding LAGs and LACP | 75 |
| | Understanding Aggregated Ethernet Interfaces and LACP for Switches | 75 |
| | Link Aggregation Group | 76 |
| | Link Aggregation Control Protocol (LACP) | 78 |
| | Configuring Aggregated Ethernet LACP (CLI Procedure) | 80 |
| | Configuring Link Aggregation | 81 |
| | Creating an Aggregated Ethernet Interface | 82 |
| | Configuring the VLAN Name and VLAN ID Number | 83 |

| | | |
|------------------|--|------------|
| | Configuring Aggregated Ethernet LACP (CLI Procedure) | 83 |
| | Configuring LACP Link Protection of Aggregated Ethernet Interfaces for | |
| | Switches | 85 |
| | Configuring LACP Link Protection for a Single Link at the Global Level | 87 |
| | Configuring LACP Link Protection for a Single Link at the Aggregated | |
| | Interface Level | 87 |
| | Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple | |
| | Links in an Aggregated Ethernet Interface | 88 |
| | Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet | |
| | Interface | 90 |
| | Verifying the Status of a LAG Interface | 90 |
| | Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging | |
| | LACP Protocol Packets | 91 |
| | Verifying the LACP Setup | 91 |
| | Verifying That LACP Packets Are Being Exchanged | 92 |
| | Troubleshooting an Aggregated Ethernet Interface | 92 |
| Part 5 | Load Balancing | |
| Chapter 5 | Understanding Load Balancing | 95 |
| | Configuring Load Balancing Based on MAC Addresses | 95 |
| Part 6 | Local Link Bias | |
| Chapter 6 | Understanding Local Link Bias | 99 |
| | Understanding Local Link Bias | 99 |
| | Configuring Local Link Bias | 101 |
| Part 7 | Redundant Trunk Groups | |
| Chapter 7 | Understanding Redundant Trunk Groups | 105 |
| | Understanding Redundant Trunk Links (Legacy RTG Configuration) | 106 |
| | Example: Configuring Redundant Trunk Links for Faster Recovery on Devices | |
| | with ELS Support | 108 |
| Part 8 | Configuration Statements and Operational Commands | |
| Chapter 8 | Ethernet OAM Link Fault Management Configuration Statements | 117 |
| | action (OAM LFM) | 118 |
| | action-profile | 119 |
| | allow-remote-loopback | 120 |
| | ethernet (Protocols OAM) | 121 |
| | event-thresholds | 127 |
| | event (LFM) | 128 |
| | frame-error | 129 |
| | frame-period | 130 |
| | frame-period-summary | 131 |
| | oam | 132 |
| | interface (OAM Link-Fault Management) | 135 |
| | link-adjacency-loss | 136 |

| | | |
|------------------|---|------------|
| | link-discovery | 137 |
| | link-down | 137 |
| | link-event-rate | 138 |
| | link-fault-management | 139 |
| | negotiation-options | 140 |
| | no-allow-link-events | 140 |
| | pdu-interval | 141 |
| | pdu-threshold | 142 |
| | remote-loopback | 143 |
| | symbol-period | 144 |
| | syslog (OAM Action) | 145 |
| | traceoptions (Individual Interfaces) | 146 |
| Chapter 9 | Interfaces Configuration Statements | 153 |
| | address | 155 |
| | alarm (chassis) | 158 |
| | auto-negotiation | 159 |
| | channel-speed | 160 |
| | configured-flow-control | 161 |
| | craft-lockout | 162 |
| | description (Interfaces) | 163 |
| | ethernet (Alarm) | 164 |
| | ethernet-switch-profile | 165 |
| | ethernet-switching | 167 |
| | ether-options | 168 |
| | eui-64 | 169 |
| | family | 170 |
| | fibre-channel (Alarm) | 174 |
| | filter | 175 |
| | flow-control | 177 |
| | fpc | 178 |
| | gratuitous-arp-reply | 179 |
| | hold-time (Physical Interface) | 180 |
| | irb (Interfaces) | 182 |
| | inet (interfaces) | 185 |
| | inet6 (interfaces) | 186 |
| | interface-mode | 187 |
| | interface-range | 189 |
| | interfaces (QFX Series) | 191 |
| | link-down | 198 |
| | link-mode | 199 |
| | link-speed | 200 |
| | loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) | 202 |
| | mac | 203 |
| | management-ethernet (Alarm) | 204 |
| | member | 205 |
| | member-range | 206 |
| | mtu | 207 |
| | no-gratuitous-arp-request | 208 |

| | | |
|-------------------|---|------------|
| | pic | 209 |
| | rx-buffers | 211 |
| | source | 212 |
| | speed (Ethernet) | 213 |
| | targeted-broadcast | 219 |
| | traceoptions (Individual Interfaces) | 220 |
| | tx-buffers | 227 |
| | unit | 229 |
| | vlan-id | 230 |
| | vlan-tagging | 231 |
| Chapter 10 | LAGs and LACP Configuration Statements | 233 |
| | aggregated-devices | 234 |
| | aggregated-ether-options | 236 |
| | chassis (QFX Series) | 239 |
| | 802.3ad | 241 |
| | device-count | 243 |
| | ethernet | 244 |
| | force-up | 245 |
| | lacp (802.3ad) | 246 |
| | lacp (Aggregated Ethernet) | 248 |
| | link-protection | 251 |
| | periodic | 252 |
| Chapter 11 | Redundant Trunk Groups Configuration Statements | 253 |
| | group (Redundant Trunk Groups) | 254 |
| | interface (Redundant Trunk Groups) | 255 |
| | preempt-cutover-timer | 256 |
| | redundant-trunk-group | 257 |
| Chapter 12 | Ethernet OAM Link Fault Management Operational Command | 259 |
| | show oam ethernet link-fault-management | 260 |
| Chapter 13 | Interfaces Operational Commands | 265 |
| | monitor interface | 266 |
| | show interfaces diagnostics optics | 279 |
| | show interfaces ge | 294 |
| | show interfaces irb | 307 |
| | show interfaces xe | 314 |
| Chapter 14 | LAGs and LACP Operational Commands | 333 |
| | show lacp interfaces | 334 |
| | show lacp statistics interfaces (View) | 339 |
| Chapter 15 | Redundant Trunk Group Operational Command | 341 |
| | show redundant-trunk-group | 342 |

List of Figures

| | | |
|------------------|--|------------|
| Part 6 | Local Link Bias | |
| Chapter 6 | Understanding Local Link Bias | 99 |
| | Figure 1: Egress Traffic Flow with Local Link Bias | 99 |
| | Figure 2: Egress Traffic Flow without Local Link Bias | 100 |
| Part 7 | Redundant Trunk Groups | |
| Chapter 7 | Understanding Redundant Trunk Groups | 105 |
| | Figure 3: Redundant Trunk Group, Link 1 Active | 107 |
| | Figure 4: Redundant Trunk Group, Link 2 Active | 107 |
| | Figure 5: Topology for Configuring the Redundant Trunk Links | 110 |

List of Tables

| | | |
|------------------|--|-------------|
| | About the Documentation | xiii |
| | Table 1: Notice Icons | xv |
| | Table 2: Text and Syntax Conventions | xvi |
| Part 2 | Interfaces | |
| Chapter 2 | Understanding Interfaces | 9 |
| | Table 3: Network Interfaces Types and Purposes for EX Series | 10 |
| | Table 4: Special Interfaces Types and Purposes for EX Series | 10 |
| | Table 5: Network Interfaces Types and Purposes for EX4600, NFX Series, QFX Series, QFabric System | 12 |
| | Table 6: Special Interfaces Types and Purposes supported on EX4600, NFX Series, QFX Series, QFabric System | 13 |
| | Table 7: Network Interfaces Types and Purposes for OCX Series | 13 |
| | Table 8: Special Interfaces Types and Purposes for OCX Series | 14 |
| | Table 9: Naming Conventions for PICs | 18 |
| | Table 10: Naming Conventions for PORTs | 19 |
| | Table 11: Configuration hierarchies for EX Series | 28 |
| | Table 12: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package | 32 |
| | Table 13: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software | 37 |
| | Table 14: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package | 40 |
| | Table 15: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software | 43 |
| | Table 16: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package | 45 |
| | Table 17: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software | 48 |
| | Table 18: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package | 51 |
| | Table 19: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software | 55 |
| Part 4 | Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP) | |
| Chapter 4 | Understanding LAGs and LACP | 75 |
| | Table 20: Maximum Interfaces per LAG and Maximum LAGs per Switch | 76 |

| | | |
|-------------------|---|------------|
| Part 7 | Redundant Trunk Groups | |
| Chapter 7 | Understanding Redundant Trunk Groups | 105 |
| | Table 21: Components of the Redundant Trunk Link Topology | 111 |
| Part 8 | Configuration Statements and Operational Commands | |
| Chapter 8 | Ethernet OAM Link Fault Management Configuration Statements | 117 |
| | Table 22: Options for traceoptions | 149 |
| Chapter 9 | Interfaces Configuration Statements | 153 |
| | Table 23: Options for speed | 217 |
| | Table 24: Options for traceoptions | 223 |
| Chapter 12 | Ethernet OAM Link Fault Management Operational Command | 259 |
| | Table 25: show oam ethernet link-fault-management Output Fields | 260 |
| Chapter 13 | Interfaces Operational Commands | 265 |
| | Table 26: Output Control Keys for the monitor interface interface-name Command | 267 |
| | Table 27: Output Control Keys for the monitor interface traffic Command | 267 |
| | Table 28: monitor interface Output Fields | 268 |
| | Table 29: show interfaces diagnostics optics Output Fields | 280 |
| | Table 30: show interfaces ge- Output Fields | 295 |
| | Table 31: show interfaces irb Output Fields | 307 |
| | Table 32: show interfaces xe Output Fields | 316 |
| Chapter 14 | LAGs and LACP Operational Commands | 333 |
| | Table 33: show lacp interfaces Output Fields | 335 |
| | Table 34: show lacp statistics interfaces Output Fields | 339 |
| Chapter 15 | Redundant Trunk Group Operational Command | 341 |
| | Table 35: show redundant-trunk-group Output Fields | 342 |

About the Documentation

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

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 <https://www.juniper.net/documentation/>.

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 <https://www.juniper.net/books>.

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see [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 |
| <i>Italic text like this</i> | <ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. | <ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i> |
| <i>Italic text like this</i> | Represents variables (options for which you substitute a value) in commands or configuration statements. | Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i> |
| Text like this | Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components. | <ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE. |
| < > (angle brackets) | Encloses optional keywords or variables. | stub <default-metric <i>metric</i> >; |
| (pipe symbol) | Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity. | broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>) |
| # (pound sign) | Indicates a comment specified on the same line as the configuration statement to which it applies. | rsvp { # Required for dynamic MPLS only |
| [] (square brackets) | Encloses a variable for which you can substitute one or more values. | community name members [community-ids] |
| Indentation and braces ({ }) | Identifies a level in the configuration hierarchy. | [edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } } |
| ;(semicolon) | Identifies a leaf statement at a configuration hierarchy level. | |

GUI Conventions


Table 2: Text and Syntax Conventions (continued)

| Convention | Description | Examples |
|--|--|---|
| Bold text like this | Represents graphical user interface (GUI) items you click or select. | <ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel. |
| > (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 so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide 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 Juniper Care or Partner Support Services 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 <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- 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 <https://support.juniper.net/support/requesting-support/>.

PART 1

Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management on page 3](#)

CHAPTER 1

Understanding Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management on page 3](#)
- [Configuring Ethernet OAM Link Fault Management on page 4](#)

Understanding Ethernet OAM Link Fault Management

Juniper Networks Junos operating system (Junos OS) for Juniper Networks allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual LAN identifier for troubleshooting.
- Work independently of the actual Ethernet transport and function over physical Ethernet ports or a virtual service such as pseudowire.
- Isolate faults over a flat (or single operator) network architecture or nested or hierarchical (or multiprovider) networks.

The following OAM LFM features are supported:

- Discovery and Link Monitoring

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.

In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The switch performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote Fault Detection

Remote fault detection uses flags and events. Flags are used to convey the following: Link Fault means a loss of signal, Dying Gasp means an unrecoverable condition such as a power failure, and Critical Event means an unspecified vendor-specific critical event. You can specify the periodic OAM PDU sending interval for fault detection. The switch uses the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote Loopback Mode

Remote loopback mode ensures link quality between the switch and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote DTE into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

Configuring Ethernet OAM Link Fault Management

Ethernet OAM link fault management (LFM) can be used for physical link-level fault detection and management. The IEEE 802.3ah LFM works across point-to-point Ethernet links either directly or through repeaters.

To configure Ethernet OAM LFM using the CLI:

1. Enable IEEE 802.3ah OAM support on an interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name
```



NOTE: You can configure Ethernet OAM LFM on aggregated interfaces.



NOTE: The remaining steps are optional. You can choose which of these features to configure for Ethernet OAM LFM on your switch.

2. Specify whether the interface or the peer initiates the discovery process by configuring the link discovery mode to **active** or **passive** (**active** = interface initiates; **passive** = peer initiates):

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name link-discovery active
```

3. Configure a periodic OAM PDU-sending interval (in milliseconds) for fault detection:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-interval interval
```

4. Specify the number of OAM PDUs that an interface can miss before the link between peers is considered down:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-threshold threshold-value
```

5. Configure event threshold values on an interface for the local errors that trigger the sending of link event TLVs:

- Set the threshold value (in seconds) for sending frame-error events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-error count
```

- Set the threshold value (in seconds) for sending frame-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period count
```

- Set the threshold value (in seconds) for sending frame-period-summary events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period-summary count
```

- Set the threshold value (in seconds) for sending symbol-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds symbol-period count
```



NOTE: You can disable the sending of link event TLVs.

To disable the sending of link event TLVs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name negotiation-options no-allow-link-events
```

6. Create an action profile to define event fault flags and thresholds to be taken when the link fault event occurs. Then apply the action profile to one or more interfaces. (You can also apply multiple action profiles to a single interface.)

- a. Name the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name
```

- b. Specify actions to be taken by the system when the link fault event occurs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name action syslog
```

```
user@switch# set action-profile profile-name action link-down
```

- c. Specify events for the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name event link-adjacency-loss
```



NOTE: For each action profile, you must specify at least one link event and one action. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all actions are executed. You can set a low threshold for a specific action such as logging the error and set a high threshold for another action such as system logging.

7. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Set the remote DTE in loopback mode (the remote DTE must support remote-loopback mode) and then enable remote loopback support for the local interface.

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name remote-loopback
```

```
user@switch# set interface interface-name negotiation-options allow-remote-loopback
```


PART 2

Interfaces

- [Understanding Interfaces on page 9](#)

CHAPTER 2

Understanding Interfaces

- [Interfaces Overview for Switches on page 9](#)
- [Understanding Interface Naming Conventions on page 14](#)
- [Understanding Interface Ranges for Switches on page 27](#)
- [Understanding Management Interfaces on page 29](#)
- [Understanding Port Ranges and System Modes on page 30](#)
- [Configuring the Interface Address on page 57](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
- [Configuring Ethernet Loopback Capability on page 65](#)
- [Monitoring Interface Status and Traffic on page 65](#)
- [Troubleshooting Network Interfaces on page 65](#)

Interfaces Overview for Switches

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

- [Network Interfaces for EX Series on page 9](#)
- [Special Interfaces for EX Series on page 10](#)
- [Network Interfaces for EX4600, NFX Series, QFX Series, QFabric System on page 12](#)
- [Special Interfaces for EX4600, NFX Series, QFX Series, QFabric System on page 13](#)
- [Network Interfaces for OCX Series on page 13](#)
- [Special Interfaces for OCX Series on page 14](#)

Network Interfaces for EX Series

Network interfaces connect to the network and carry network traffic. [Table 3 on page 10](#) lists the types of network interfaces supported on EX Series switches.

Table 3: Network Interfaces Types and Purposes for EX Series

| Type | Purpose |
|--------------------------------------|---|
| Aggregated Ethernet interfaces | <p>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 <i>link aggregation group (LAG)</i> or <i>bundle</i>. These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.</p> <p>See “Understanding Aggregated Ethernet Interfaces and LACP for Switches” on page 75.</p> |
| LAN access interfaces | <p>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.</p> |
| Power over Ethernet (PoE) interfaces | <p>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.</p> <p>See Understanding PoE on EX Series Switches.</p> |
| Trunk interfaces | <p>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.</p> |

Special Interfaces for EX Series

Table 4 on page 10 lists the types of special interfaces supported on EX Series switches.

Table 4: Special Interfaces Types and Purposes for EX Series

| Type | Purpose |
|----------------------|---|
| Console port | <p>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 since 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.</p> |
| Loopback | <p>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.</p> |
| Management interface | <p>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.)</p> <p>See “Understanding Management Interfaces” on page 29.</p> |

Table 4: Special Interfaces Types and Purposes for EX Series (continued)

| Type | Purpose |
|--|--|
| Integrated Routing and Bridging (IRB) Interface or Routed VLAN Interface (RVI) | <p>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.</p> <p>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.</p> <p>See <i>Understanding Integrated Routing and Bridging</i>.</p> |
| Virtual Chassis port (VCP) interfaces | <p>Virtual Chassis ports (VCPs) are used to interconnect switches in a Virtual Chassis:</p> <ul style="list-style-type: none"> 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 <i>Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port</i>. EX4200 and EX4500 switches—Each EX4200 switch or each EX4500 switch with a Virtual Chassis module installed has two dedicated VCPs on its rear panel. These ports can be used to interconnect up to ten EX4200 switches in an EX4200 Virtual Chassis, up to ten EX4500 switches in an EX4500 Virtual Chassis, and up to ten switches in a mixed EX4200 and EX4500 Virtual Chassis. When you power on switches that are interconnected in this manner, the software automatically configures the VCP interfaces for the dedicated ports that have been interconnected. These VCP interfaces are not configurable or modifiable. See <i>Understanding the High-Speed Interconnection of the Dedicated Virtual Chassis Ports Connecting EX4200, EX4500, and EX4550 Member Switches</i>. <p>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 <i>Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port</i>.</p> <ul style="list-style-type: none"> EX4300 switches—All QSFP+ ports are configured as VCPs by default. See <i>Understanding EX Series Virtual Chassis</i>. <p>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 <i>Setting an Uplink Port on an EX Series or QFX Series Switch as a Virtual Chassis Port</i>.</p> <ul style="list-style-type: none"> 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. <p>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.</p> |
| Virtual management Ethernet (VME) interface | <p>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 <i>Understanding Global Management of a Virtual Chassis</i>.</p> <p>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.</p> |

Network Interfaces for EX4600, NFX Series, QFX Series, QFabric System

Network interfaces connect to the network and carry network traffic. [Table 5 on page 12](#) lists the types of network interfaces supported.

Table 5: Network Interfaces Types and Purposes for EX4600, NFX Series, QFX Series, QFabric System

| Type | Purpose |
|---|---|
| Aggregated Ethernet interfaces | Group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth. |
| Channelized Interfaces | <p>Depending on the device and software package, 40-Gbps QSFP+ ports can be configured to operate as the following types of interfaces:</p> <ul style="list-style-type: none"> 10-Gigabit Ethernet interfaces (<i>xe</i>) 40-Gigabit Ethernet interfaces (<i>et</i> and <i>xle</i>) 40-Gigabit data plane uplink interfaces (<i>fte</i>) <p>When an <i>et</i> port is channelized to four <i>xe</i> ports, a colon is used to signify the four separate channels. For example, on a QFX3500 standalone switch with port 2 on PIC 1 configured as four 10-Gigabit Ethernet ports, the interface names are <i>xe-0/1/2:0</i>, <i>xe-0/1/2:1</i>, <i>xe-0/1/2:2</i>, and <i>xe-0/1/2:3</i></p> <p>NOTE: You cannot configure channelized interfaces to operate as Virtual Chassis ports.</p> |
| Ethernet Interfaces | Configure Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet interfaces to connect to other servers, storage, and switches. You can configure 40-Gigabit data plane uplink ports to connect a Node device to an Interconnect devices as well as for Virtual Chassis ports (VCPs). |
| Fibre Channel interfaces | <p>Use Fibre Channel interfaces to connect the switch to a Fibre Channel over Ethernet (FCoE) forwarder or a Fibre Channel switch in a storage area network (SAN). You can configure Fibre Channel interfaces only on ports 0 through 5 and 42 through 47 on QFX3500 devices. Fibre Channel interfaces do not forward Ethernet traffic.</p> <p>See <i>Overview of Fibre Channel</i>.</p> |
| LAN access interfaces | Use these interfaces to connect to other servers, storage, and switches. When you power on a QFX Series product and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports. |
| Multichassis aggregated Ethernet (MC-AE) interfaces | Group a LAG on one standalone switch with a LAG on another standalone switch to create a MC-AE. The MC-AE provides load balancing and redundancy across the two standalone switches. |
| Tagged-access mode interfaces | Use tagged-access interfaces to connect a switch to an access layer device. Tagged-access interfaces can accept VLAN-tagged packets from multiple VLANs. |
| Trunk interfaces | Use trunk interfaces to connect to other switches or routers. To use a port for this type of connection, you must explicitly configure the port interface for trunk mode. The interfaces from the switches or routers must also be configured for trunk mode. In this mode, the interface can be in multiple VLANs and accept tagged packets from multiple devices. Trunk interfaces typically connect to other switches and to routers on the LAN. |

Table 5: Network Interfaces Types and Purposes for EX4600, NFX Series, QFX Series, QFabric System (continued)

| Type | Purpose |
|------------------------------|--|
| Virtual Chassis ports (VCPs) | You can use Virtual Chassis ports to send and receive Virtual Chassis Control Protocol (VCCP) traffic, and to create, monitor, and maintain the Virtual Chassis. On QFX3500, QFX3600, QFX5100, QFX5110, QFX5200, and EX4600 standalone switches, you can configure 40-Gigabit Ethernet QSFP+ uplink ports (non-channelized) or fixed SFP+ 10-Gigabit Ethernet ports as VCPs by issuing the request virtual-chassis-vc-port-set CLI command. QFX5110 switches also support configuring 100-Gigabit QSFP28 ports as VCPs. |

Special Interfaces for EX4600, NFX Series, QFX Series, QFabric System

Table 6 on page 13 lists the types of special interfaces supported.

Table 6: Special Interfaces Types and Purposes supported on EX4600, NFX Series, QFX Series, QFabric System

| Type | Purpose |
|---|---|
| Console port | Each device has a serial console port, labeled CON or CONSOLE , for connecting tty-type terminals to the switch. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch. |
| Loopback interface | A software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch. |
| Management interface | <p>The management Ethernet interface provides an out-of-band method for connecting to a standalone switch and QFabric system.</p> <p>NOTE: On OCX Series switches, the em0 management interface always has the status up in show command outputs, even if the physical port is empty. The me0 interface is a virtual interface between Junos and the host operating system, therefore its status is independent from the status of the physical port.</p> |
| Routed VLAN interfaces (RVI and IRB interfaces) | <p>Layer 3 routed VLAN interfaces (called RVI in the original CLI, and called IRB in Enhanced Layer 2 Software) route traffic from one broadcast domain to another and perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network.</p> <p>The RVI or IRB functions as a logical router, eliminating the need for having both a switch and a router. The RVI or IRB must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it.</p> |

Network Interfaces for OCX Series

Network interfaces connect to the network and carry network traffic. Table 7 on page 13 lists the types of network interfaces supported.

Table 7: Network Interfaces Types and Purposes for OCX Series

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

Special Interfaces for OCX Series

Table 8 on page 14 lists the types of special interfaces supported.

Table 8: Special Interfaces Types and Purposes for OCX Series

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

Related Documentation

- [EX2200 Switches Hardware Overview](#)
- [EX3200 System 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](#)
- [Understanding Layer 3 Logical Interfaces on page 69](#)
- [Understanding Layer 3 Subinterfaces](#)

Understanding Interface Naming Conventions

The EX Series, QFX Series, NFX Series, OCX1100, QFabric System, and EX4600 devices use a naming convention for defining the interfaces that are similar to that of other platforms running under Juniper Networks Junos OS. This topic provides brief information about the naming conventions used for interfaces on the QFX Series and on EX4600 switches.

For detailed information on interface naming like physical part, logical part, and channel part of the interfaces, see [Interface Naming Overview](#).

This topic describes:

- [Physical Part of an Interface Name for EX Series on page 15](#)
- [Logical Part of an Interface Name for EX Series on page 16](#)
- [Wildcard Characters in Interface Names for EX Series on page 17](#)
- [Physical Part of an Interface Name for QFX series, NFX Series, EX4600, QFabric System on page 17](#)
- [Logical Part of an Interface Name on a Switch Running QFabric Software Package for QFX series, NFX Series, EX4600, QFabric System on page 25](#)
- [Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software for QFX series, NFX Series, EX4600, QFabric System on page 25](#)
- [Wildcard Characters in Interface Names for QFX series, NFX Series, EX4600, QFabric System on page 26](#)
- [Physical Part of an Interface Name for OCX1100 on page 26](#)
- [Wildcard Characters in Interface Names for OCX1100 on page 26](#)

Physical Part of an Interface Name for EX Series

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 for EX Series

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 for EX Series

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 (" ").

Physical Part of an Interface Name for QFX series, NFX Series, EX4600, QFabric System

Interfaces in Junos OS are specified as follows:

device-name:type-fpc/pic/port

The convention is as follows (and platform support depends on the Junos OS release in your installation):

- *device-name*—(QFabric systems only) The *device-name* is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.
- *type*—The QFX Series and EX4600 device interfaces use the following media types:
 - **fc**—Fibre Channel interface
 - **ge**—Gigabit Ethernet interface
 - **xe**—10-Gigabit Ethernet interface
 - **sxe**—10-Gigabit Service interface. **sxe** is an internal interface and user must not configure this interface. It supports L2 and L3 configurations like VLANs and IP address.
 - **xle**—40-Gigabit Ethernet interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
 - **et**—40-Gigabit Ethernet interface (QFX3500, QFX3600, QFX5100, QFX5200, QFX10000, and EX4600 switches running Enhanced Layer 2 Software)
 - **et**—100-Gigabit Ethernet interface (QFX5200 and QFX10000 switches running Enhanced Layer 2 Software)
 - **fte**—40-Gigabit data plane uplink interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
 - **me**—Management interface
 - **em**—Management interface on QFX5100 and EX4600 switches.
- *fpc*—Flexible PIC Concentrator. QFX Series interfaces use the following convention for the FPC number in interface names:
 - On QFX3500, QFX3600, QFX5100 devices running a QFabric software package, and QFX10002 switches, the FPC number is always 0.

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

- On QFX3500, QFX3600, QFX5100, QFX5200, EX4600, QFX10002, QFX10008, and QFX10016 switches running Enhanced Layer 2 Software, the member ID of a member in a Virtual Chassis determines the FPC number.



NOTE: Every member in a Virtual Chassis must have a unique member ID, otherwise the Virtual Chassis will not be created.

- On standalone QFX5100, EX4600, and QFX10002 switches, the FPC number is always 0.
- pic*—QFX Series and EX4600 device interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:

Table 9: Naming Conventions for PICs

| Device with Software Package | Convention |
|--|---|
| QFX3500 switch with QFabric software package | PIC 0 can support 48 ports, PIC 1 can support 16 10-Gigabit Ethernet ports, and PIC 2 can support 4 40-Gigabit Ethernet ports. |
| QFX3500 switch with Enhanced Layer 2 software | PIC 0 can support 48 ports, and PIC 1 can support 16 10-Gigabit Ethernet ports, and 4 40-Gigabit Ethernet ports. |
| QFX3500 Node device with a QFabric software package | PIC 0 can support 48 ports and PIC 1 can support four 40-Gigabit data plane uplink ports. |
| QFX3600 switch with a QFabric software package | PIC 0 can support 64 10-Gigabit Ethernet ports, and PIC 1 can support 16 40-Gigabit Ethernet ports. |
| QFX3600 switch with Enhanced Layer 2 software | PIC 0 can support 64 10-Gigabit Ethernet ports and can also support 16 40-Gigabit Ethernet ports. |
| QFX3600 Node device running a QFabric software package | PIC 0 can support 56 10-Gigabit Ethernet ports, and PIC 1 can support 8 40-Gigabit data plane uplink ports, and up to 14 40-Gigabit Ethernet ports. |
| QFX5100-48S switch with Enhanced Layer 2 software | PIC 0 provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces. |
| EX4600 device with Enhanced Layer 2 software | PIC 0 provides 4 40-Gbps QSFP+ ports and 24 10-Gigabit Ethernet interfaces. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provide 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 10-Gbps SFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each. |
| QFX5100-48S switch with a QFabric software package | PIC 1 provides six 40-Gbps QSFP+ ports, and PIC 0 provides 48 10-Gigabit Ethernet interfaces. |
| QFX5100-24Q switch with Enhanced Layer 2 software | PIC 0 provides 24 40-Gbps QSFP+ ports. PIC 1 and PIC 2 can each contain a QFX-EM-4Q expansion module, and each expansion module provides 4 40-Gbps QSFP+ ports |
| QFX5100-96S switch with Enhanced Layer 2 software | PIC 0 provides 96 10-Gigabit Ethernet interfaces and 8 40-Gbps QSFP+ ports . |

Table 9: Naming Conventions for PICs (continued)

| Device with Software Package | Convention |
|--|--|
| QFX5110-48S switch with Enhanced Layer 2 software | PIC 0 can support 48 10-Gigabit Ethernet ports labeled 0 through 47, and 4 QSFP28 ports labeled 48 through 51. Ports 0 through 47 support either 1-Gbps small form-factor pluggable (SFP) or 10-Gbps small form-factor pluggable plus (SFP+) transceivers. You can also use SFP+ DAC cables and 10-Gbps active optical cables (AOC) in any access port. The default 100-Gigabit Ethernet ports can be configured as 40-Gigabit Ethernet, and in this configuration can either operate as dedicated 40-Gigabit Ethernet ports or can be channelized to 4 independent 10-Gigabit Ethernet ports using copper or fiber breakout cables. |
| QFX5200-32C switch with Enhanced Layer 2 software | PIC 0 provides 32 QSFP28 ports. The 100-Gigabit Ethernet ports can be channelized to two 50-Gigabit Ethernet or four 25-Gigabit Ethernet ports. The default 100-Gigabit Ethernet ports can be configured as 40-Gigabit Ethernet and operate as 40-Gigabit Ethernet or be channelized to four 10-Gigabit Ethernet ports. |
| QFX10002-36Q switch with Enhanced Layer 2 software | PIC 0 provides 144 10-Gigabit Ethernet interfaces, and 36 40-Gbps QSFP+ ports, and 12 100-Gigabit Ethernet interfaces. |
| QFX10002-72Q switch with Enhanced Layer 2 software | PIC 0 provides 288 10-Gigabit Ethernet interfaces, and 72 40-Gbps QSFP+ ports, and 24 100-Gigabit Ethernet interfaces. |
| QFX10008 switch with Enhanced Layer 2 software | PIC 0 provides one-thousand, one-hundred fifty two 10-Gigabit Ethernet interfaces, two-hundred eighty-eight 40-Gbps QSFP+ ports, or two-hundred forty 100-Gigabit Ethernet interfaces. |
| QFX10016 switch with Enhanced Layer 2 software | PIC 0 provides two-thousand, three-hundred and four 10-Gigabit Ethernet interfaces, five-hundred seventy-six 40-Gbps QSFP+ ports, or four-hundred eighty 100-Gigabit Ethernet interfaces. |

- *port*—Interfaces use the following convention for port numbers:

Table 10: Naming Conventions for PORTs

| Device with Software Package | Convention |
|--|--|
| QFX3500 switch with a QFabric software package | <p>There are 48 network access ports (10-Gigabit Ethernet) labeled 0 through 47 on PIC 0 and, 16 network access ports labeled 0 through 15 on PIC 1, and four 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 2. You can use the QSFP+ ports to connect the Node device to Interconnect devices.</p> <p>By default, the 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet ports. You can use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. Optionally, you can choose to configure the QSFP+ ports as 40-Gigabit Ethernet ports (see <i>Configuring the QSFP+ Port Type on QFX3500 Standalone Switches</i>).</p> |
| QFX3500 switch with Enhanced Layer 2 software | There are 48 network access ports labeled 0 through 47 on PIC 0 and 4 40-Gbps QSFP+ ports labeled Q0 through Q3 on PIC 1. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports. |

Table 10: Naming Conventions for PORTs (continued)

| Device with Software Package | Convention |
|---|--|
| QFX3600 switch with a QFabric software package | <p>There are 64 network access ports (10-Gigabit Ethernet) labeled Q0 through Q15 on PIC 0, and there are 16 network access ports (40-Gigabit Ethernet) labeled Q0 through Q15 on PIC 1.</p> <p>By default, all the QSFP+ ports are configured to operate as 40-Gigabit Ethernet ports. Optionally, you can choose to configure the QSFP+ ports as 10-Gigabit Ethernet ports (see <i>Configuring the Port Type on QFX3600 Standalone Switches</i>) and use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.</p> |
| QFX3600 Node device with a QFabric software package | <p>PIC 0 can support up to 56 10-Gigabit Ethernet ports labeled Q2 through Q15, and PIC 1 can support up to 8 40-Gigabit data plane uplink ports labeled Q0 through Q7, and up to 14 40-Gigabit Ethernet ports labeled Q2 through Q15.</p> <p>On a QFX3600 Node device, by default, four 40-Gbps QSFP+ ports (labeled Q0 through Q3) are configured for uplink connections between your Node device and your Interconnect devices, and twelve 40-Gbps QSFP+ ports (labeled Q4 through Q15) use QSFP+ to four SFP+ copper breakout cables to support up to 48 10-Gigabit Ethernet ports for connections to either endpoint systems (such as servers and storage devices) or external networks. Optionally, you can choose to configure the first eight ports (Q0 through Q7) for uplink connections between your Node device and your Interconnect devices, and ports Q2 through Q15 for 10-Gigabit Ethernet or 40-Gigabit Ethernet connections to either endpoint systems or external networks (see <i>Configuring the Port Type on QFX3600 Node Devices</i>).</p> |
| QFX3600 switch with Enhanced Layer 2 software | PIC 0 can support 64 network access ports (10-Gigabit Ethernet ports) labeled Q0 through Q15 and 16 40-Gigabit Ethernet ports labeled Q0 through Q15. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports. |
| QFX5100-48S switch with Enhanced Layer 2 software | PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47 and 6 40-Gbps QSFP+ ports labeled 48 through 53. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports. |
| EX4600 switch with Enhanced Layer 2 software | PIC 0 can support 24 network access ports (10-Gigabit Ethernet ports) labeled 0 through 23 and 4 40-Gbps QSFP+ ports labeled 24 through 27. There are two expansion bays (PIC 1 and PIC 2), and you can insert QFX-EM-4Q expansion modules and EX4600-EM-8F expansion modules. The QFX-EM-4Q expansion module provide 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 10-Gbps SFP+ ports. You can insert any combination of expansion modules. For example, you can insert two EX4600-EM-8F expansion modules, two QFX-EM-4Q expansion modules, or one of each. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports. |
| QFX5100-48S switch with a QFabric software package | PIC 0 can support 48 network access ports (10-Gigabit Ethernet ports) labeled 0 through 47, and PIC 1 can support 6 40-Gbps QSFP+ ports labeled 0 through 5. See <i>Configuring the QSFP+ Port Type on QFX5100 Devices</i> for information on how to configure the port mode of 40-Gbps QSFP+ ports. |

Table 10: Naming Conventions for PORTs (continued)

| Device with Software Package | Convention |
|---|--|
| QFX5100-24Q switch with Enhanced Layer 2 software | <p>PIC 0 can support 24 40-Gbps QSFP+ ports labeled 0 through 23. PIC 1 and PIC 2 each support 4 40-Gbps QSFP+ port, for a total of eight 40-Gbps QSFP+ ports. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>NOTE: You cannot channelize the 40-Gbps QSFP+ ports provided in the two QFX-EM-4Q expansion modules. Also, even though there is a total of 128 physical ports, only 104 logical ports can be channelized.</p> <p>You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See <i>Configuring the System Mode</i> for information on how to configure the system mode.</p> |
| QFX5100-96S switch with Enhanced Layer 2 software | <p>PIC 0 can support 96 10-Gigabit Ethernet ports labeled 0 through 95, and 8 40-Gbps QSFP+ ports labeled 96 through 103. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>NOTE: You can only channelize the 40-Gbps QSFP+ ports provided in ports 96 and 100, because only 104 logical ports can be channelized.</p> <p>You can configure different system modes to achieve varying levels of port density on the QFX5100-24Q and QFX5100-96S switches. Depending on the system mode you configure, there are restrictions on which ports you can channelize. If you channelize ports that are restricted, the configuration is ignored. See <i>Configuring the System Mode</i> for information on how to configure the system mode.</p> |
| QFX5110-48S switch with Enhanced Layer 2 software | <p>PIC 0 can support 48 10-Gigabit Ethernet ports labeled 0 through 47, and 4 QSFP28 ports labeled 48 through 51. These data ports (0 through 47) support either 1-Gbps small form-factor pluggable (SFP) or 10-Gbps small form-factor pluggable plus (SFP+) transceivers. You can also use SFP+ DAC cables and 10-Gbps active optical cables (AOC) in any access port. The default 100-Gigabit Ethernet ports can be configured as 40-Gigabit Ethernet, and in this configuration can either operate as dedicated 40-Gigabit Ethernet ports or can be channelized to 4 independent 10-Gigabit Ethernet ports using copper or fiber breakout cables.</p> |
| QFX5200-32C switch with Enhanced Layer 2 software | <p>There is support for both quad small-form-factor pluggable (QSFP+) and 28-Gbps QSFP+ (QSFP28) transceivers in the 32 QSFP28 sockets. The QSFP28 ports are configured as 100-Gigabit Ethernet ports by default, but can also be configured to speeds of 50, 40, 25, or 10 Gigabit Ethernet.</p> <p>The 100 Gigabit Ethernet ports can be channelized using breakout cables either to 2 independent downstream 50 Gigabit Ethernet or to 4 independent 25 Gigabit Ethernet ports. The default 100 Gigabit Ethernet ports can also be configured as 40 Gigabit Ethernet and in this configuration can either operate as dedicated 40 Gigabit Ethernet ports or can be channelized to 4 independent 10 Gigabit Ethernet ports using breakout cables. See <i>Channelizing Interfaces on QFX5200-32C Switches</i> for information on how to configure and channelize the interfaces.</p> <p>NOTE: Autochannelization is not supported.</p> |

Table 10: Naming Conventions for PORTs (continued)

| Device with Software Package | Convention |
|--|--|
| QFX10002-36Q switch with Enhanced Layer 2 software | <p>There are 36 quad small-form factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 36 ports, 12 ports are QSFP28 capable, which are dual speed 40- or 100-Gigabit Ethernet optical transceivers.</p> <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using 28-Gbps QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. • 10-Gigabit Ethernet using breakout cables. When configured for channelization, a breakout cable converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports. <p>Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 12 QSFP28 ports support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers • 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports support:</p> <ul style="list-style-type: none"> • 40-Gigabit Ethernet QSFP+ transceivers • Access ports |
| QFX10002-72Q switch with Enhanced Layer 2 software | <p>There are 72 quad small-form factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 72 ports, 24 ports are QSFP28 capable, which are dual speed 40- or 100-Gigabit Ethernet optical transceivers.</p> <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using 28-Gbps QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. • 10-Gigabit Ethernet using breakout cables. When configured for channelization, a breakout cable converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports. <p>Any of the 72 ports 0 through 71 can be configured as either uplink or access ports. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 24 QSFP28 ports support:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers <p>Each of the 72 QSFP+ ports support:</p> <ul style="list-style-type: none"> • 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports support:</p> <ul style="list-style-type: none"> • 40-Gigabit Ethernet QSFP+ transceivers • Access ports • Uplink ports |

Table 10: Naming Conventions for PORTs (continued)

| Device with Software Package | Convention |
|--|---|
| On a QFX10008 switch with Enhanced Layer 2 software, there are two line cards available: | QFX10000-36Q, a 36-port 40-Gigabit Ethernet quad small form-factor pluggable plus transceiver (QSFP+) or 12-port 100GbE QSFP28 line card |
| QFX10008 with Line Card QFX10000-36Q (ELS) | <p>The QFX10000-36Q line cards supports</p> <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet. 40-Gigabit Ethernet using QSFP+ optical transceivers. 10-Gigabit Ethernet using breakout cabling and attached optical transceivers. When configured for channelization, the system converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports. <p>Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 12 QSFP28 ports supports:</p> <ul style="list-style-type: none"> 100-Gigabit Ethernet QSFP28 transceivers 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 12 QSFP28 ports supports:</p> <ul style="list-style-type: none"> 100-Gigabit Ethernet QSFP28 transceivers 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports support:</p> <ul style="list-style-type: none"> 40-Gigabit Ethernet QSFP+ transceivers Access ports Uplink ports |
| QFX10008 with Line Card QFX10000-30C and QFX10000-30C-M (ELS) | <p>QFX10000-30C and QFX10000-30C-M, a 30-port 100-Gigabit or 40-Gigabit Ethernet QSFP28 line card</p> <ul style="list-style-type: none"> The QFX10000-30C and QFX10000-30C-M line cards support: Thirty 28-Gbps QSFP+ Pluggable Solution (QSFP28) cages that support either 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The QFX10000-30C and QFX10000-30C-M ports auto detect the type of transceiver installed and set the configuration to the appropriate speed. <p>Each QSFP28 socket can be configured to support:</p> <ul style="list-style-type: none"> 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet. 40-Gigabit Ethernet using QSFP+ optical transceivers. <p>See <i>Channelizing Interfaces Overview</i> for information on how to configure and channelize the 40-Gbps QSFP+ ports.</p> <p>Each of the 30 QSFP28 ports supports:</p> <ul style="list-style-type: none"> 100-Gigabit Ethernet QSFP28 transceivers 40-Gigabit Ethernet QSFP+ transceivers Access ports Uplink ports |

Table 10: Naming Conventions for PORTs (continued)

| Device with Software Package | Convention |
|---|---|
| On a QFX10016 switch running Enhanced Layer 2 software, there are 16 slots, which you can populate with two types line cards: QFX10016 with Line Card QFX10000-36Q (ELS) | <ul style="list-style-type: none"> QFX10000-36Q, a 36-port 40-Gigabit Ethernet quad small form-factor pluggable plus transceiver (QSFP+) or 12-port 100GbE QSFP28 line card <p>The QFX10000-36Q line card consists of 36 quad small form-factor pluggable plus (QSFP+) ports that support 40-Gigabit Ethernet optical transceivers. Out of these 36 ports, 12 ports are QSFP28 capable. The QSFP+ ports are dual speed and can support either 40-Gigabit or 100-Gigabit Ethernet optical transceivers. The line card can support 10-Gigabit Ethernet by channelizing the 40-Gigabit ports. Channelization is supported on fiber break-out cable using standard structured cabling techniques.</p> <p>With 100-Gigabit Ethernet using QSFP28 optical transceivers, when a QSFP28 transceiver is inserted into the ports marked with a fine black line underneath the socket and the port is configured for 100-Gigabit Ethernet, the two adjacent ports are disabled and the QSFP28 socket is enabled for 100-Gigabit Ethernet.</p> <p>You can use 40-Gigabit Ethernet using QSFP+ optical transceivers.</p> <p>With 10-Gigabit Ethernet using breakout cabling and attached optical transceivers, when configured for channelization, the system converts the 40-Gigabit Ethernet port into 4 independent 10-Gigabit Ethernet ports.</p> <p>Any of the 36 ports 0 through 35 can be configured as either uplink or access ports.</p> <p>Each of the 12 QSFP28 ports supports:</p> <ul style="list-style-type: none"> 100-Gigabit Ethernet QSFP28 transceivers 40-Gigabit Ethernet QSFP+ transceivers <p>Each of the 36 QSFP+ ports supports:</p> <ul style="list-style-type: none"> 40-Gigabit Ethernet QSFP+ transceivers Access ports <p>You can use 40-Gigabit Ethernet QSFP+ transceivers in any downstream port.</p> Uplink ports <p>You can configure all the QSFP+ ports as uplinks.</p> <p>Every second and sixth port in a 6XQSFP cage on a QFX10000-36Q supports 100-Gigabit Ethernet using QSFP28 transceivers. These 100-Gigabit Ethernet ports work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet, but are recognized as 40-Gigabit Ethernet by default. When a 40-Gigabit Ethernet transceiver is inserted into a 100-Gigabit Ethernet port, the port recognizes the 40-Gigabit Ethernet port speed. When a 100-Gigabit Ethernet transceiver is inserted into the port and enabled in the CLI, the port recognizes the 100-Gigabit Ethernet speed and disables two adjacent 40-Gigabit Ethernet ports. You can also use an 100-Gigabit Ethernet transceiver and run it at 40-Gigabit Ethernet by using the CLI to set the port speed to 40-Gigabit Ethernet.</p> <p>The 40-Gigabit Ethernet ports can operate independently, be channelized into four 10-Gigabit Ethernet ports, or bundled with the next two consecutive ports and channelized into twelve 10-Gigabit Ethernet ports as a port range. Only the first and fourth port in each 6XQSFP cage are available to channelize a port range. The port range must be configured using the <code>set chassis fpc 0 pic 0 port 1 channel-speed 10g</code> command.</p> |

Table 10: Naming Conventions for PORTs (continued)

| Device with Software Package | Convention |
|---|--|
| QFX10016 with Line Card QFX10000-30C and QFX10000-30C-M (ELS) | <p>The QFX10000-30C and QFX10000-30C-M line cards consist of thirty 28-Gbps QSFP+ Pluggable Solution (QSFP28) cages that support either 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The QFX10000-30C and QFX10000-30C-M ports auto detect the type of transceiver installed and set the configuration to the appropriate speed.</p> <p>Each QSFP28 socket supports:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet using QSFP28 optical transceivers. When a QSFP28 transceiver is inserted into any of the ports, the QSFP28 socket is enabled for 100-Gigabit Ethernet. • 40-Gigabit Ethernet using QSFP+ optical transceivers. When a QSFP+ transceiver is inserted into any of the ports, the QSFP+ socket is enabled for 40-Gigabit. <p>Any of the 30 ports 0 through 29 can be configured as either uplink or access ports, and of the 30 QSFP28 ports supports:</p> <ul style="list-style-type: none"> • 100-Gigabit Ethernet QSFP28 transceivers • 40-Gigabit Ethernet QSFP+ transceivers |

Logical Part of an Interface Name on a Switch Running QFabric Software Package for QFX series, NFX Series, EX4600, QFabric System

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: *device-name* (QFabric systems only): *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 |
|-------------------------|-------|-----------------|-----------|
| node-device1:xe-0/0/1.0 | down | remote-analyzer | unblocked |
| node-device1:xe-0/0/2.0 | down | default | unblocked |
| node-device1:xe-0/0/3.0 | down | default | unblocked |

When you configure aggregated Ethernet interfaces, you configure a logical interface, which is called a *bundle* or a *LAG*. Each LAG can include up to eight Ethernet interfaces, depending on the switch model.

Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software for QFX series, NFX Series, EX4600, QFabric System

Channelizing enables you to configure four 10-Gigabit Ethernet interfaces from a 40-Gigabit Ethernet QSFP+ interface. By default, a 40-Gigabit Ethernet QSFP+ interface is named *et-fpc/pic/port*. The resulting 10-Gigabit Ethernet interfaces appear in the following format: *xe-fpc/pic/port:channel*, where channel can be a value of 0 through 3.

For example, if an *et* interface named **et-0/0/3** is channelized to four 10-Gigabit Ethernet interfaces, the resulting 10-Gigabit Ethernet interface names will be **xe-0/0/3:0**, **xe-0/0/3:1**, **xe-0/0/3:2**, and **xe-0/0/3:3**:

| Interface | Admin | Link | Proto | Local | Remote |
|------------|-------|------|-------|-------|--------|
| xe-0/0/3:0 | up | down | | | |
| xe-0/0/3:1 | up | down | | | |
| xe-0/0/3:2 | up | down | | | |
| xe-0/0/3:3 | up | down | | | |

Wildcard Characters in Interface Names for QFX series, NFX Series, EX4600, QFabric System

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 (" ").

Physical Part of an Interface Name for OCX1100

Interfaces in Junos OS are specified as follows:

type-fpc/pic/port

The convention is as follows:

- *type*—The OCX Series device interfaces use the following media types:
 - **xe**—10-Gigabit Ethernet interface
 - **et**—40-Gigabit Ethernet interface
 - **em**—Management interface
- *fpc*—Flexible PIC Concentrator. OCX Series interfaces use the following convention for the FPC number in interface names:
 - On standalone OCX Series switches, the FPC number is always **0**.
The FPC number indicates the slot number of the line card that contains the physical interface.
- *pic*—The OCX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
 - PIC **0** provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces.
- *port*—Interfaces use the following convention for port numbers:
 - PIC **0** can support 48 network access ports (10-Gigabit Ethernet ports) labeled 1 through 48 and 6 40-Gbps QSFP+ ports labeled 49 through 54.

Wildcard Characters in Interface Names for OCX1100

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the *interface-name* option to specify groups of interface names without having to type

each name individually. You must enclose all wildcard characters except the asterisk (*) in quotation marks (" ").

Related Documentation

- [Interfaces Overview for Switches on page 9](#)
- [Channelizing Interfaces Overview](#)
- [Understanding Management Interfaces on page 29](#)
- [Understanding Port Ranges and System Modes on page 30](#)
- [Configuring the System Mode](#)
- [Configuring Gigabit Ethernet Interfaces \(CLI Procedure\)](#)
- [Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)
- [Rear Panel of a QFX3500 Device](#)
- [Front Panel of a QFX3600 Device](#)

Understanding Interface Ranges for Switches

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

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



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

Starting in Junos OS Release 14.1X53-D15 and later, the common configurations defined in the interface range will not be overridden but appended to the local configuration. In Junos OS Releases prior to 14.1X53-D15, 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:

[Table 11 on page 28](#) lists the configuration hierarchies for the EX Series, NFX, OCX, QFX Series, and QFabric Series.

Table 11: Configuration hierarchies for EX Series

| Configuration Hierarchies for EX Series | Configuration Hierarchies for EX4600, NFX, QFX Series, and QFabric Systems | Configuration Hierarchies for EX Series with ELS |
|--|--|---|
| <ul style="list-style-type: none"> • ethernet-switching-options analyzer <i>name</i> input egress interface • ethernet-switching-options analyzer <i>name</i> input ingress interface • ethernet-switching-options analyzer output interface • ethernet-switching-options bpdu-block interface • ethernet-switching-options interfaces • ethernet-switching-options redundant-trunk-group group-name interface • ethernet-switching-options secure-access-port interface • ethernet-switching-options voip interface • poe interface • protocols dot1x authentication interface • protocols gvrp interface • protocols igmp interface • protocols igmp-snooping vlan <i>vlan-name</i> interface • protocols isis interface • protocols link-management peer lmp-control-channel interface • protocols link-management te-link <i>name</i> interface • protocols lldp interface • protocols lldp-med interface • protocols mpls interface • protocols mstp interface • protocols mstp msti-<i>id</i> interface • protocols mstp msti-<i>id</i> vlan <i>vlan-id</i> interface • protocols oam ethernet link-fault-management interface • protocols ospf area • protocols pim interface • protocols rip group <i>group-name</i> neighbor • protocols ripng group <i>group-name</i> neighbor • protocols router-advertisement interface • protocols router-discovery interface • protocols rsvp interface • protocols sflow interfaces • protocols stp interface • protocols vstp vlan <i>vlan-id</i> interface • vlans <i>vlan-name</i> interface | <ul style="list-style-type: none"> • protocols isis interface • protocols sflow interfaces <p>NOTE: These statements are not supported on OCX Series switches.</p> | <ul style="list-style-type: none"> • forwarding-options analyzer <i>name</i> input egress interface • forwarding-options analyzer <i>name</i> 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 <i>name</i> lmp-control-channel • protocols link-management te-link <i>name</i> interface • protocols lldp interface • protocols lldp-med interface • protocols mstp interface • protocols oam ethernet link-fault-management interface • protocols ospf area <i>area-id</i> interface • protocols pim interface • protocols router-advertisement interface • protocols router-discovery interface • protocols rsvp interface • protocols sflow interfaces • protocols vstp vlan <i>vlan-id</i> interface • switch-options redundant-trunk-group <i>group-name</i> interface • switch-options voip interface <p>For ELS details, see <i>Using the Enhanced Layer 2 Software CLI</i>.</p> |

- Related Documentation**
- *Configuring Interface Ranges*
 - *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
 - *Configuring Aggregated Ethernet Links (CLI Procedure)*
 - *Configuring a Layer 3 Subinterface (CLI Procedure)*
 - *interface-range*
 - [Configuring Link Aggregation on page 81](#)
 - [Configuring a Layer 3 Logical Interface on page 70](#)
 - *Junos OS Network Interfaces Library for Routing Devices*

Understanding Management Interfaces

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



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

Management interface port ranges vary based on device type (and platform support depends on the Junos OS release in your installation):

- QFX3500 devices:

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

- QFX3600 devices:

There are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0s** and **C1s**). On a QFX3600 standalone switch, the RJ-45 management interfaces and SFP management interfaces correspond to **me0** and **me1**. On a QFX3600 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**. Each pair of management interfaces correspond to one Ethernet interface—for example, both RJ-45 management interfaces (labeled **C0** and **C0s**) can correspond to **me0**, and both SFP management interfaces (labeled **C1** and **C1s**) can correspond to **me1**. By default, both RJ-45

management interfaces are active. If you insert an SFP interface into the SFP management port (**COS**, for example), the SFP interface would become the active management interface, and the corresponding RJ-45 management interface (**C0**) is disabled.



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

- On QFX5100, QFX5200, and EX4600 switches, there is one RJ-45 management interface (labeled **C0**) and one SFP management interface (labeled **C1**), and they correspond to em0 and em1. You can use both management interfaces simultaneously.
- On QFX10002 and QFX10008 switches, there is one RJ-45 management interface (labeled **MGMT**) and one SFP management interface (labeled **MGMT**), and they correspond to em0 and em1. Although the CLI permits you to configure two management Ethernet interfaces within the same subnet, only one interface is usable and supported.

- On OCX Series switches:

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

- QFabric system:

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

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

Understanding Port Ranges and System Modes

QFX Series devices and EX4600 switches can support different port ranges depending on the device, media type of the interface, the software that is running on the device, and the system mode.

This topic describes:

- [Port Ranges for Different Media Types on page 31](#)
- [Supported System Modes on page 54](#)

Port Ranges for Different Media Types

The following media types support the following port ranges:

- On a QFX3500 device:
 - The valid port range for a Fibre Channel (fc) interface is **0** through **5** and **42** through **47** on PIC **0**, with a total of 12 available Fibre Channel ports.



NOTE: Fibre Channel ports are not supported on QFX3500, QFX3600, and QFX5100 switches running Enhanced Layer 2 software.

- The valid port range for a Gigabit Ethernet (ge) interface is **6** through **41** on PIC **0** because the ports between **0** and **5** and **42** and **47** are reserved as Fibre Channel ports. The total number of available Gigabit Ethernet ports is 36, because 12 of the remaining 48 ports are reserved for Fibre Channel and 10-Gigabit Ethernet interfaces. Fibre Channel ports cannot be configured as Gigabit Ethernet ports.
- The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **47** on PIC **0**. The valid port range for a 10-Gigabit Ethernet (xe) interface is **0** through **15** on PIC **1**. The total number of available 10-Gigabit Ethernet ports is 64.
- The valid port range for a 40-Gigabit data plane uplink interface is **0** through **3** on PIC **1**
- The valid port range for a 40-Gigabit Ethernet interface is **0** through **3** on PIC **2**. There are four available ports.
- On a QFX3600 Node device:
 - The valid port range for a 10-Gigabit Ethernet interface is **8** through **63** on PIC **0**. There are 56 available ports.
 - The valid port range for a 40-Gigabit Ethernet interface is **2** through **15** on PIC **1**. There are 14 available ports.
 - The valid port range for a 40-Gigabit data plane uplink interface is **0** through **7** on PIC **1**. There are eight available ports.

See [Table 14 on page 40](#) for physical port to logical port mappings.

- On a QFX3600 switch running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **63** on PIC **0**. There are 64 available ports.
 - The valid port range for a 40-Gigabit Ethernet interface is **0** through **15** on PIC **0**. There are 16 available ports.

See [Table 15 on page 43](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running Enhanced Layer 2 Software:

- The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC **0**. There are 48 available ports. When you channelize the 6 40-Gbps QSFP+ ports on **0** through **5** on PIC **1**, there are 72 available ports.



NOTE: On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

- The valid port range for a 40-Gbps QSFP+ port is **0** through **5** on PIC **1**. There are six available ports.

See [Table 17 on page 48](#) for physical port to logical port mappings.

- On EX4600 switches running Enhanced Layer 2 Software:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **23** on PIC **0**. There are 24 available ports. When you channelize the 4 40-Gbps QSFP+ ports on **24** through **27** on PIC **0**. There are 40 available ports.

See [Table 17 on page 48](#) for physical port to logical port mappings.

- On QFX5100-48S and QFX5100-48T switches running a QFabric software package:
 - The valid port range for a 10-Gigabit Ethernet interface is **0** through **47** on PIC **0**. There are 48 available ports.
 - The valid port range for a 40-Gbps QSFP+ port is **0** through **5** on PIC **1**. There are six available ports.



NOTE: On PIC 1, ports 0 and 1 are reserved for fte ports. You cannot convert these fte ports to xe or xle ports.

See [Table 18 on page 51](#) for physical port to logical port mappings.

- For QFX5100-24Q and QFX5100-96S switches running Enhanced Layer 2 Software, see [Table 19 on page 55](#) for physical port to logical port mappings for different system modes.

Table 12: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package

| Port Number | Fibre Channel Interfaces (On PIC 0) | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 2) |
|-------------|--|---|--|--|--|
| 0 | fc-0/0/0 | Not supported on this port | xe-0/0/0 | Not supported on this port | Not supported on this port |
| 1 | fc-0/0/1 | Not supported on this port | xe-0/0/1 | Not supported on this port | Not supported on this port |

Table 12: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

| Port Number | Fibre Channel Interfaces (On PIC 0) | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 2) |
|-------------|--|---|--|--|--|
| 2 | fc-0/0/2 | Not supported on this port | xe-0/0/2 | Not supported on this port | Not supported on this port |
| 3 | fc-0/0/3 | Not supported on this port | xe-0/0/3 | Not supported on this port | Not supported on this port |
| 4 | fc-0/0/4 | Not supported on this port | xe-0/0/4 | Not supported on this port | Not supported on this port |
| 5 | fc-0/0/5 | Not supported on this port | xe-0/0/5 | Not supported on this port | Not supported on this port |
| 6 | Not supported on this port | ge-0/0/6 | xe-0/0/6 | Not supported on this port | Not supported on this port |
| 7 | Not supported on this port | ge-0/0/7 | xe-0/0/7 | Not supported on this port | Not supported on this port |
| 8 | Not supported on this port | ge-0/0/8 | xe-0/0/8 | Not supported on this port | Not supported on this port |
| 9 | Not supported on this port | ge-0/0/9 | xe-0/0/9 | Not supported on this port | Not supported on this port |
| 10 | Not supported on this port | ge-0/0/10 | xe-0/0/10 | Not supported on this port | Not supported on this port |
| 11 | Not supported on this port | ge-0/0/11 | xe-0/0/11 | Not supported on this port | Not supported on this port |
| 12 | Not supported on this port | ge-0/0/12 | xe-0/0/12 | Not supported on this port | Not supported on this port |
| 13 | Not supported on this port | ge-0/0/13 | xe-0/0/13 | Not supported on this port | Not supported on this port |
| 14 | Not supported on this port | ge-0/0/14 | xe-0/0/14 | Not supported on this port | Not supported on this port |
| 15 | Not supported on this port | ge-0/0/15 | xe-0/0/15 | Not supported on this port | Not supported on this port |
| 16 | Not supported on this port | ge-0/0/16 | xe-0/0/16 | Not supported on this port | Not supported on this port |

Table 12: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

| Port Number | Fibre Channel Interfaces (On PIC 0) | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 2) |
|-------------|--|---|--|--|--|
| 17 | Not supported on this port | ge-0/0/17 | xe-0/0/17 | Not supported on this port | Not supported on this port |
| 18 | Not supported on this port | ge-0/0/18 | xe-0/0/18 | Not supported on this port | Not supported on this port |
| 19 | Not supported on this port | ge-0/0/19 | xe-0/0/19 | Not supported on this port | Not supported on this port |
| 20 | Not supported on this port | ge-0/0/20 | xe-0/0/20 | Not supported on this port | Not supported on this port |
| 21 | Not supported on this port | ge-0/0/21 | xe-0/0/21 | Not supported on this port | Not supported on this port |
| 22 | Not supported on this port | ge-0/0/22 | xe-0/0/22 | Not supported on this port | Not supported on this port |
| 23 | Not supported on this port | ge-0/0/23 | xe-0/0/23 | Not supported on this port | Not supported on this port |
| 24 | Not supported on this port | ge-0/0/24 | xe-0/0/24 | Not supported on this port | Not supported on this port |
| 25 | Not supported on this port | ge-0/0/25 | xe-0/0/25 | Not supported on this port | Not supported on this port |
| 26 | Not supported on this port | ge-0/0/26 | xe-0/0/26 | Not supported on this port | Not supported on this port |
| 27 | Not supported on this port | ge-0/0/27 | xe-0/0/27 | Not supported on this port | Not supported on this port |
| 28 | Not supported on this port | ge-0/0/28 | xe-0/0/28 | Not supported on this port | Not supported on this port |
| 29 | Not supported on this port | ge-0/0/29 | xe-0/0/29 | Not supported on this port | Not supported on this port |
| 30 | Not supported on this port | ge-0/0/30 | xe-0/0/30 | Not supported on this port | Not supported on this port |
| 31 | Not supported on this port | ge-0/0/31 | xe-0/0/31 | Not supported on this port | Not supported on this port |

Table 12: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

| Port Number | Fibre Channel Interfaces (On PIC 0) | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 2) |
|-------------|--|---|--|--|--|
| 32 | Not supported on this port | ge-0/0/32 | xe-0/0/32 | Not supported on this port | Not supported on this port |
| 33 | Not supported on this port | ge-0/0/33 | xe-0/0/33 | Not supported on this port | Not supported on this port |
| 34 | Not supported on this port | ge-0/0/34 | xe-0/0/34 | Not supported on this port | Not supported on this port |
| 35 | Not supported on this port | ge-0/0/35 | xe-0/0/35 | Not supported on this port | Not supported on this port |
| 36 | Not supported on this port | ge-0/0/36 | xe-0/0/36 | Not supported on this port | Not supported on this port |
| 37 | Not supported on this port | ge-0/0/37 | xe-0/0/37 | Not supported on this port | Not supported on this port |
| 38 | Not supported on this port | ge-0/0/38 | xe-0/0/38 | Not supported on this port | Not supported on this port |
| 39 | Not supported on this port | ge-0/0/39 | xe-0/0/39 | Not supported on this port | Not supported on this port |
| 40 | Not supported on this port | ge-0/0/40 | xe-0/0/40 | Not supported on this port | Not supported on this port |
| 41 | Not supported on this port | ge-0/0/41 | xe-0/0/41 | Not supported on this port | Not supported on this port |
| 42 | fc-0/0/42 | Not supported on this port | xe-0/0/42 | Not supported on this port | Not supported on this port |
| 43 | fc-0/0/43 | Not supported on this port | xe-0/0/43 | Not supported on this port | Not supported on this port |
| 44 | fc-0/0/44 | Not supported on this port | xe-0/0/44 | Not supported on this port | Not supported on this port |
| 45 | fc-0/0/45 | Not supported on this port | xe-0/0/45 | Not supported on this port | Not supported on this port |
| 46 | fc-0/0/46 | Not supported on this port | xe-0/0/46 | Not supported on this port | Not supported on this port |

Table 12: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

| Port Number | Fibre Channel Interfaces (On PIC 0) | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 2) |
|-------------|--|---|--|--|--|
| 47 | fc-0/0/47 | Not supported on this port | xe-0/0/47 | Not supported on this port | Not supported on this port |
| Q0 | Not supported on this port | Not supported on this port | xe-0/1/0 xe-0/1/1 xe-0/1/2 xe-0/1/3 NOTE: Supported on QFX3500 standalone switch only. | fte-0/1/0 | xle-0/2/0 |
| Q1 | Not supported on this port | Not supported on this port | xe-0/1/4 xe-0/1/5 xe-0/1/6 xe-0/1/7 NOTE: Supported on QFX3500 standalone switch only. | fte-0/1/1 | xle-0/2/1 |
| Q2 | Not supported on this port | Not supported on this port | xe-0/1/8 xe-0/1/9 xe-0/1/10 xe-0/1/11 NOTE: Supported on QFX3500 standalone switch only. | fte-0/1/2 | xle-0/2/2 |

Table 12: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package (continued)

| Port Number | Fibre Channel Interfaces (On PIC 0) | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 2) |
|-------------|--|---|--|--|--|
| Q3 | Not supported on this port | Not supported on this port | xe-0/1/12 xe-0/1/13 xe-0/1/14 xe-0/1/15 NOTE: Supported on QFX3500 standalone switch only. | fte-0/1/3 | xle-0/2/3 |

Table 13: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software

| Port Number | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|---|--|--|
| 0 | Not supported on this port | xe-0/0/0 | Not supported on this port |
| 1 | Not supported on this port | xe-0/0/1 | Not supported on this port |
| 2 | Not supported on this port | xe-0/0/2 | Not supported on this port |
| 3 | Not supported on this port | xe-0/0/3 | Not supported on this port |
| 4 | Not supported on this port | xe-0/0/4 | Not supported on this port |
| 5 | Not supported on this port | xe-0/0/5 | Not supported on this port |
| 6 | ge-0/0/6 | xe-0/0/6 | Not supported on this port |
| 7 | ge-0/0/7 | xe-0/0/7 | Not supported on this port |
| 8 | ge-0/0/8 | xe-0/0/8 | Not supported on this port |
| 9 | ge-0/0/9 | xe-0/0/9 | Not supported on this port |
| 10 | ge-0/0/10 | xe-0/0/10 | Not supported on this port |
| 11 | ge-0/0/11 | xe-0/0/11 | Not supported on this port |
| 12 | ge-0/0/12 | xe-0/0/12 | Not supported on this port |

Table 13: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (continued)

| Port Number | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|---|--|--|
| 13 | ge-0/0/13 | xe-0/0/13 | Not supported on this port |
| 14 | ge-0/0/14 | xe-0/0/14 | Not supported on this port |
| 15 | ge-0/0/15 | xe-0/0/15 | Not supported on this port |
| 16 | ge-0/0/16 | xe-0/0/16 | Not supported on this port |
| 17 | ge-0/0/17 | xe-0/0/17 | Not supported on this port |
| 18 | ge-0/0/18 | xe-0/0/18 | Not supported on this port |
| 19 | ge-0/0/19 | xe-0/0/19 | Not supported on this port |
| 20 | ge-0/0/20 | xe-0/0/20 | Not supported on this port |
| 21 | ge-0/0/21 | xe-0/0/21 | Not supported on this port |
| 22 | ge-0/0/22 | xe-0/0/22 | Not supported on this port |
| 23 | ge-0/0/23 | xe-0/0/23 | Not supported on this port |
| 24 | ge-0/0/24 | xe-0/0/24 | Not supported on this port |
| 25 | ge-0/0/25 | xe-0/0/25 | Not supported on this port |
| 26 | ge-0/0/26 | xe-0/0/26 | Not supported on this port |
| 27 | ge-0/0/27 | xe-0/0/27 | Not supported on this port |
| 28 | ge-0/0/28 | xe-0/0/28 | Not supported on this port |
| 29 | ge-0/0/29 | xe-0/0/29 | Not supported on this port |
| 30 | ge-0/0/30 | xe-0/0/30 | Not supported on this port |
| 31 | ge-0/0/31 | xe-0/0/31 | Not supported on this port |
| 32 | ge-0/0/32 | xe-0/0/32 | Not supported on this port |
| 33 | ge-0/0/33 | xe-0/0/33 | Not supported on this port |
| 34 | ge-0/0/34 | xe-0/0/34 | Not supported on this port |
| 35 | ge-0/0/35 | xe-0/0/35 | Not supported on this port |

Table 13: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (continued)

| Port Number | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|---|--|--|
| 36 | ge-0/0/36 | xe-0/0/36 | Not supported on this port |
| 37 | ge-0/0/37 | xe-0/0/37 | Not supported on this port |
| 38 | ge-0/0/38 | xe-0/0/38 | Not supported on this port |
| 39 | ge-0/0/39 | xe-0/0/39 | Not supported on this port |
| 40 | ge-0/0/40 | xe-0/0/40 | Not supported on this port |
| 41 | ge-0/0/41 | xe-0/0/41 | Not supported on this port |
| 42 | Not supported on this port | xe-0/0/42 | Not supported on this port |
| 43 | Not supported on this port | xe-0/0/43 | Not supported on this port |
| 44 | Not supported on this port | xe-0/0/44 | Not supported on this port |
| 45 | Not supported on this port | xe-0/0/45 | Not supported on this port |
| 46 | Not supported on this port | xe-0/0/46 | Not supported on this port |
| 47 | Not supported on this port | xe-0/0/47 | Not supported on this port |
| Q0 | Not supported on this port | xe-0/1/0:0 xe-0/1/0:1 xe-0/1/0:2 xe-0/1/0:3 | et-0/1/0 |
| Q1 | Not supported on this port | xe-0/1/1:0 xe-0/1/1:1 xe-0/1/1:2 xe-0/1/1:3 | et-0/1/1 |
| Q2 | Not supported on this port | xe-0/1/2:0 xe-0/1/2:1 xe-0/1/2:2 xe-0/1/2:3 | et-0/1/2 |

Table 13: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software (continued)

| Port Number | Gigabit Ethernet Interfaces (On PIC 0) | 10-Gigabit Ethernet Interfaces (On PIC 0 and 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|---|--|--|
| Q3 | Not supported on this port | xe-0/1/3:0 xe-0/1/3:1 xe-0/1/3:2 xe-0/1/3:3 | et-0/1/3 |

Table 14: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|--|--|
| Q0 | xe-0/0/0 xe-0/0/1 xe-0/0/2 xe-0/0/3 | xle-0/1/0 |
| Q1 | xe-0/0/4 xe-0/0/5 xe-0/0/6 xe-0/0/7 | xle-0/1/1 |
| Q2 | xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11 | xle-0/1/2 |
| Q3 | xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15 | xle-0/1/3 |

Table 14: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|--|--|
| Q4 | xe-0/0/16 | xle-0/1/4 |
| | xe-0/0/17 | |
| | xe-0/0/18 | |
| | xe-0/0/19 | |
| Q5 | xe-0/0/20 | xle-0/1/5 |
| | xe-0/0/21 | |
| | xe-0/0/22 | |
| | xe-0/0/23 | |
| Q6 | xe-0/0/24 | xle-0/1/6 |
| | xe-0/0/25 | |
| | xe-0/0/26 | |
| | xe-0/0/27 | |
| Q7 | xe-0/0/28 | xle-0/1/7 |
| | xe-0/0/29 | |
| | xe-0/0/30 | |
| | xe-0/0/31 | |
| Q8 | xe-0/0/32 | xle-0/1/8 |
| | xe-0/0/33 | |
| | xe-0/0/34 | |
| | xe-0/0/35 | |
| Q9 | xe-0/0/36 | xle-0/1/9 |
| | xe-0/0/37 | |
| | xe-0/0/38 | |
| | xe-0/0/39 | |

Table 14: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|--|--|
| Q10 | xe-0/0/40 | xle-0/1/10 |
| | xe-0/0/41 | |
| | xe-0/0/42 | |
| | xe-0/0/43 | |
| Q11 | xe-0/0/44 | xle-0/1/11 |
| | xe-0/0/45 | |
| | xe-0/0/46 | |
| | xe-0/0/47 | |
| Q12 | xe-0/0/48 | xle-0/1/12 |
| | xe-0/0/49 | |
| | xe-0/0/50 | |
| | xe-0/0/51 | |
| Q13 | xe-0/0/52 | xle-0/1/13 |
| | xe-0/0/53 | |
| | xe-0/0/54 | |
| | xe-0/0/55 | |
| Q14 | xe-0/0/56 | xle-0/1/14 |
| | xe-0/0/57 | |
| | xe-0/0/58 | |
| | xe-0/0/59 | |
| Q15 | xe-0/0/60 | xle-0/1/15 |
| | xe-0/0/61 | |
| | xe-0/0/62 | |
| | xe-0/0/63 | |

Table 15: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 0) |
|-------------|--|--|
| Q0 | xe-0/0/0:0 | et-0/0/0 |
| | xe-0/0/0:1 | |
| | xe-0/0/0:2 | |
| | xe-0/0/0:3 | |
| Q1 | xe-0/0/1:0 | et-0/0/1 |
| | xe-0/0/1:1 | |
| | xe-0/0/1:2 | |
| | xe-0/0/1:3 | |
| Q2 | xe-0/0/2:0 | et-0/0/2 |
| | xe-0/0/2:1 | |
| | xe-0/0/2:2 | |
| | xe-0/0/2:3 | |
| Q3 | xe-0/0/3:0 | et-0/0/3 |
| | xe-0/0/3:1 | |
| | xe-0/0/3:2 | |
| | xe-0/0/3:3 | |
| Q4 | xe-0/0/4:0 | et-0/0/4 |
| | xe-0/0/4:1 | |
| | xe-0/0/4:2 | |
| | xe-0/0/4:3 | |
| Q5 | xe-0/0/5:0 | et-0/0/5 |
| | xe-0/0/5:1 | |
| | xe-0/0/5:2 | |
| | xe-0/0/5:3 | |

Table 15: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 0) |
|-------------|--|--|
| Q6 | xe-0/0/6:0 | et-0/0/6 |
| | xe-0/0/6:1 | |
| | xe-0/0/6:2 | |
| | xe-0/0/6:3 | |
| Q7 | xe-0/0/7:0 | et-0/0/7 |
| | xe-0/0/7:1 | |
| | xe-0/0/7:2 | |
| | xe-0/0/7:3 | |
| Q8 | xe-0/0/8:0 | et-0/0/8 |
| | xe-0/0/8:1 | |
| | xe-0/0/8:2 | |
| | xe-0/0/8:3 | |
| Q9 | xe-0/0/9:0 | et-0/0/9 |
| | xe-0/0/9:1 | |
| | xe-0/0/9:2 | |
| | xe-0/0/9:3 | |
| Q10 | xe-0/0/10:0 | et-0/0/10 |
| | xe-0/0/10:1 | |
| | xe-0/0/10:2 | |
| | xe-0/0/10:3 | |
| Q11 | xe-0/0/11:0 | et-0/0/11 |
| | xe-0/0/11:1 | |
| | xe-0/0/11:2 | |
| | xe-0/0/11:3 | |

Table 15: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 0) |
|-------------|--|--|
| Q12 | xe-0/0/12:0 xe-0/0/12:1 xe-0/0/12:2 xe-0/0/12:3 | et-0/0/12 |
| Q13 | xe-0/0/13:0 xe-0/0/13:1 xe-0/0/13:2 xe-0/0/13:3 | et-0/0/13 |
| Q14 | xe-0/0/14:0 xe-0/0/14:1 xe-0/0/14:2 xe-0/0/14:3 | et-0/0/14 |
| Q15 | xe-0/0/15:0 xe-0/0/15:1 xe-0/0/15:2 xe-0/0/15:3 | et-0/0/15 |

Table 16: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|--|---|--|
| Q0 | Not supported on this port | fte-0/1/0 | xle-0/1/0 |
| Q1 | Not supported on this port | fte-0/1/1 | xle-0/1/1 |
| Q2 | xe-0/0/8 xe-0/0/9 xe-0/0/10 xe-0/0/11 | fte-0/1/2 | xle-0/1/2 |

Table 16: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|--|--|---|
| Q3 | xe-0/0/12 | fte-0/1/3 | xle-0/1/3 |
| | xe-0/0/13 | | |
| | xe-0/0/14 | | |
| | xe-0/0/15 | | |
| Q4 | xe-0/0/16 | fte-0/1/4 | xle-0/1/4 |
| | xe-0/0/17 | | |
| | xe-0/0/18 | | |
| | xe-0/0/19 | | |
| Q5 | xe-0/0/20 | fte-0/1/5 | xle-0/1/5 |
| | xe-0/0/21 | | |
| | xe-0/0/22 | | |
| | xe-0/0/23 | | |
| Q6 | xe-0/0/24 | fte-0/1/6 | xle-0/1/6 |
| | xe-0/0/25 | | |
| | xe-0/0/26 | | |
| | xe-0/0/27 | | |
| Q7 | xe-0/0/28 | fte-0/1/7 | xle-0/1/7 |
| | xe-0/0/29 | | |
| | xe-0/0/30 | | |
| | xe-0/0/31 | | |
| Q8 | xe-0/0/32 | Not supported on this port | xle-0/1/8 |
| | xe-0/0/33 | | |
| | xe-0/0/34 | | |
| | xe-0/0/35 | | |

Table 16: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|--|--|---|
| Q9 | xe-0/0/36 | Not supported on this port | xle-0/1/9 |
| | xe-0/0/37 | | |
| | xe-0/0/38 | | |
| | xe-0/0/39 | | |
| Q10 | xe-0/0/40 | Not supported on this port | xle-0/1/10 |
| | xe-0/0/41 | | |
| | xe-0/0/42 | | |
| | xe-0/0/43 | | |
| Q11 | xe-0/0/44 | Not supported on this port | xle-0/1/11 |
| | xe-0/0/45 | | |
| | xe-0/0/46 | | |
| | xe-0/0/47 | | |
| Q12 | xe-0/0/48 | Not supported on this port | xle-0/1/12 |
| | xe-0/0/49 | | |
| | xe-0/0/50 | | |
| | xe-0/0/51 | | |
| Q13 | xe-0/0/52 | Not supported on this port | xle-0/1/13 |
| | xe-0/0/53 | | |
| | xe-0/0/54 | | |
| | xe-0/0/55 | | |
| Q14 | xe-0/0/56 | Not supported on this port | xle-0/1/14 |
| | xe-0/0/57 | | |
| | xe-0/0/58 | | |
| | xe-0/0/59 | | |

Table 16: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) | 40-Gigabit Ethernet Interfaces (On PIC 1) |
|-------------|--|--|---|
| Q15 | xe-0/0/60 xe-0/0/61 xe-0/0/62 xe-0/0/63 | Not supported on this port | xle-0/1/15 |

Table 17: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 0) |
|-------------|--|--|
| 0 | xe-0/0/0 | Not supported on this port |
| 1 | xe-0/0/1 | Not supported on this port |
| 2 | xe-0/0/2 | Not supported on this port |
| 3 | xe-0/0/3 | Not supported on this port |
| 4 | xe-0/0/4 | Not supported on this port |
| 5 | xe-0/0/5 | Not supported on this port |
| 6 | xe-0/0/6 | Not supported on this port |
| 7 | xe-0/0/7 | Not supported on this port |
| 8 | xe-0/0/8 | Not supported on this port |
| 9 | xe-0/0/9 | Not supported on this port |
| 10 | xe-0/0/10 | Not supported on this port |
| 11 | xe-0/0/11 | Not supported on this port |
| 12 | xe-0/0/12 | Not supported on this port |
| 13 | xe-0/0/13 | Not supported on this port |
| 14 | xe-0/0/14 | Not supported on this port |
| 15 | xe-0/0/15 | Not supported on this port |

Table 17: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 0) |
|-------------|--|--|
| 16 | xe-0/0/16 | Not supported on this port |
| 17 | xe-0/0/17 | Not supported on this port |
| 18 | xe-0/0/18 | Not supported on this port |
| 19 | xe-0/0/19 | Not supported on this port |
| 20 | xe-0/0/20 | Not supported on this port |
| 21 | xe-0/0/21 | Not supported on this port |
| 22 | xe-0/0/22 | Not supported on this port |
| 23 | xe-0/0/23 | Not supported on this port |
| 24 | xe-0/0/24 | Not supported on this port |
| 25 | xe-0/0/25 | Not supported on this port |
| 26 | xe-0/0/26 | Not supported on this port |
| 27 | xe-0/0/27 | Not supported on this port |
| 28 | xe-0/0/28 | Not supported on this port |
| 29 | xe-0/0/29 | Not supported on this port |
| 30 | xe-0/0/30 | Not supported on this port |
| 31 | xe-0/0/31 | Not supported on this port |
| 32 | xe-0/0/32 | Not supported on this port |
| 33 | xe-0/0/33 | Not supported on this port |
| 34 | xe-0/0/34 | Not supported on this port |
| 35 | xe-0/0/35 | Not supported on this port |
| 36 | xe-0/0/36 | Not supported on this port |
| 37 | xe-0/0/37 | Not supported on this port |

Table 17: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 0) |
|-------------|--|--|
| 38 | xe-0/0/38 | Not supported on this port |
| 39 | xe-0/0/39 | Not supported on this port |
| 40 | xe-0/0/40 | Not supported on this port |
| 41 | xe-0/0/41 | Not supported on this port |
| 42 | xe-0/0/42 | Not supported on this port |
| 43 | xe-0/0/43 | Not supported on this port |
| 44 | xe-0/0/44 | Not supported on this port |
| 45 | xe-0/0/45 | Not supported on this port |
| 46 | xe-0/0/46 | Not supported on this port |
| 47 | xe-0/0/47 | Not supported on this port |
| 48 | xe-0/0/48:0 xe-0/0/48:1 xe-0/0/48:2 xe-0/0/48:3 | et-0/0/48 |
| 49 | xe-0/0/49:0 xe-0/0/49:1 xe-0/0/49:2 xe-0/0/49:3 | et-0/0/49 |
| 50 | xe-0/0/50:0 xe-0/0/50:1 xe-0/0/50:2 xe-0/0/50:3 | et-0/0/50 |

Table 17: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 0) |
|-------------|--|--|
| 51 | xe-0/0/51:0 xe-0/0/51:1 xe-0/0/51:2 xe-0/0/51:3 | et-0/0/51 |
| 52 | xe-0/0/52:0 xe-0/0/52:1 xe-0/0/52:2 xe-0/0/52:3 | et-0/0/52 |
| 53 | xe-0/0/53:0 xe-0/0/53:1 xe-0/0/53:2 xe-0/0/53:3 | et-0/0/53 |

Table 18: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) |
|-------------|---|--|--|
| 0 | xe-0/0/0 | Not supported on this port | Not supported on this port |
| 1 | xe-0/0/1 | Not supported on this port | Not supported on this port |
| 2 | xe-0/0/2 | Not supported on this port | Not supported on this port |
| 3 | xe-0/0/3 | Not supported on this port | Not supported on this port |
| 4 | xe-0/0/4 | Not supported on this port | Not supported on this port |
| 5 | xe-0/0/5 | Not supported on this port | Not supported on this port |
| 6 | xe-0/0/6 | Not supported on this port | Not supported on this port |
| 7 | xe-0/0/7 | Not supported on this port | Not supported on this port |
| 8 | xe-0/0/8 | Not supported on this port | Not supported on this port |

Table 18: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) |
|-------------|--|---|---|
| 9 | xe-0/0/9 | Not supported on this port | Not supported on this port |
| 10 | xe-0/0/10 | Not supported on this port | Not supported on this port |
| 11 | xe-0/0/11 | Not supported on this port | Not supported on this port |
| 12 | xe-0/0/12 | Not supported on this port | Not supported on this port |
| 13 | xe-0/0/13 | Not supported on this port | Not supported on this port |
| 14 | xe-0/0/14 | Not supported on this port | Not supported on this port |
| 15 | xe-0/0/15 | Not supported on this port | Not supported on this port |
| 16 | xe-0/0/16 | Not supported on this port | Not supported on this port |
| 17 | xe-0/0/17 | Not supported on this port | Not supported on this port |
| 18 | xe-0/0/18 | Not supported on this port | Not supported on this port |
| 19 | xe-0/0/19 | Not supported on this port | Not supported on this port |
| 20 | xe-0/0/20 | Not supported on this port | Not supported on this port |
| 21 | xe-0/0/21 | Not supported on this port | Not supported on this port |
| 22 | xe-0/0/22 | Not supported on this port | Not supported on this port |
| 23 | xe-0/0/23 | Not supported on this port | Not supported on this port |
| 24 | xe-0/0/24 | Not supported on this port | Not supported on this port |
| 25 | xe-0/0/25 | Not supported on this port | Not supported on this port |
| 26 | xe-0/0/26 | Not supported on this port | Not supported on this port |
| 27 | xe-0/0/27 | Not supported on this port | Not supported on this port |
| 28 | xe-0/0/28 | Not supported on this port | Not supported on this port |
| 29 | xe-0/0/29 | Not supported on this port | Not supported on this port |
| 30 | xe-0/0/30 | Not supported on this port | Not supported on this port |

Table 18: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) |
|-------------|--|---|---|
| 31 | xe-0/0/31 | Not supported on this port | Not supported on this port |
| 32 | xe-0/0/32 | Not supported on this port | Not supported on this port |
| 33 | xe-0/0/33 | Not supported on this port | Not supported on this port |
| 34 | xe-0/0/34 | Not supported on this port | Not supported on this port |
| 35 | xe-0/0/35 | Not supported on this port | Not supported on this port |
| 36 | xe-0/0/36 | Not supported on this port | Not supported on this port |
| 37 | xe-0/0/37 | Not supported on this port | Not supported on this port |
| 38 | xe-0/0/38 | Not supported on this port | Not supported on this port |
| 39 | xe-0/0/39 | Not supported on this port | Not supported on this port |
| 40 | xe-0/0/40 | Not supported on this port | Not supported on this port |
| 41 | xe-0/0/41 | Not supported on this port | Not supported on this port |
| 42 | xe-0/0/42 | Not supported on this port | Not supported on this port |
| 43 | xe-0/0/43 | Not supported on this port | Not supported on this port |
| 44 | xe-0/0/44 | Not supported on this port | Not supported on this port |
| 45 | xe-0/0/45 | Not supported on this port | Not supported on this port |
| 46 | xe-0/0/46 | Not supported on this port | Not supported on this port |
| 47 | xe-0/0/47 | Not supported on this port | Not supported on this port |
| 48 | Not supported on this port | Not supported on this PIC | fte-0/1/0 NOTE: This interface is a fixed fte interface and cannot be changed to xle. |
| 49 | Not supported on this port | Not supported on this PIC | fte-0/1/1 NOTE: This interface is a fixed fte interface and cannot be changed to xle. |

Table 18: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package (continued)

| Port Number | 10-Gigabit Ethernet Interfaces (On PIC 0) | 40-Gigabit Ethernet Interfaces (On PIC 1) | 40-Gigabit Data Plane Uplink Interfaces (On PIC 1) |
|-------------|--|---|---|
| 50 | Not supported on this port | xle-0/1/2 | fte-0/1/2 <i>NOTE:</i> By default, this interface is an fte interface but can be configured as an xle interface. |
| 51 | Not supported on this port | xle-0/1/3 | fte-0/1/3 <i>NOTE:</i> By default, this interface is an fte interface but can be configured as an xle interface. |
| 52 | Not supported on this port | xle-0/1/4 <i>NOTE:</i> By default, this interface is an xle interface but can be configured as an fte interface. | fte-0/1/4 |
| 53 | Not supported on this port | xle-0/1/5 <i>NOTE:</i> By default, this interface is an xle interface but can be configured as an fte interface. | fte-0/1/5 |

Supported System Modes



NOTE: There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches depending on the system mode you configure. If you try to channelize ports that are restricted, the configuration is ignored.

The following system modes are available on the QFX5100-24Q switch:

- Default mode
- Mode-104-port
- Flexi-PIC mode
- Non-oversubscribed mode

See [Table 19 on page 55](#) for more information regarding the supported system modes for your switch.

The following system modes are available on the QFX5100-96S switch:

- Default-mode
- Non-oversubscribed mode

See [Table 19 on page 55](#) for more information regarding the supported system modes for your switch.

Table 19: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software

| | Default-mode | Mode-104port | Flexi-pic-mode | Non-oversubscribed-mode |
|-----------------------------|--|--|---|---|
| QFX5100-48S and QFX5100-48T | Not supported | Not supported | Not supported | Not supported |
| QFX5100-24Q | Supported You do not need to configure the switch to be in this mode. On PIC 0, you can channelize all 24 40-Gbps QSFP+ ports. On PIC 1 and PIC 2, the 40-Gbps QSFP+ ports in the expansion modules are supported but cannot be channelized. In this mode, you can have one of two port combinations: 32 40-Gbps QSFP+ ports, or 96 10-Gigabit Ethernet ports plus 8 40-Gbps QSFP+ ports. | Supported On PIC 0, all 24 40-Gbps QSFP+ ports are channelized by default, which provides 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in an expansion module on PIC 1 are supported. On PIC 1, ports 0 and 2 are channelized by default, and ports 1 and 3 are disabled. If 40-Gbps QSFP+ ports contained in an expansion module are detected on PIC 2, they are ignored. | Supported On PIC 0, the first four ports (ports 0 through 3) cannot be channelized. 40-Gbps QSFP+ ports contained in expansion modules on PIC 1 and PIC 2 are supported but cannot be channelized. | Supported All 24 40-Gbps QSFP+ ports on PIC 0 can be channelized to 96 10-Gigabit Ethernet ports. 40-Gbps QSFP+ ports contained in the expansion modules on PIC 1 and PIC 2 are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode. |
| QFX5100-96S | Supported You do not need to configure the switch to be in this mode. On PIC 0, all 96 10-Gigabit Ethernet ports are supported. You can only channelize the 40-Gbps QSFP+ interfaces to 10-Gigabit Ethernet interfaces on ports 96 and 100. When you channelize the interfaces on ports 96 and 100, ports 97, 98, 99, 101, 102 and 103 are disabled. | Not supported | Not supported | Supported On PIC 0, all 96 10-Gigabit Ethernet ports are supported. However, the eight 40-Gbps QSFP+ ports are not supported and cannot be channelized. There is no packet loss for packets of any size in this mode. |

- Related Documentation**
- *Rear Panel of a QFX3500 Device*
 - *Front Panel of a QFX3600 Device*

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 **vppls** 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:bbb:...:zzzz/nn*** statement at the **[edit interfaces]** hierarchy level.

```
[edit interfaces ]
```

```
user@host# set interface-name unit logical-unit-number family inet6 address  
aaaa:bbb:...: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 itself 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 Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches

Devices include a factory default configuration that:

- Enables all 10-Gigabit Ethernet network interfaces on the switch
- Sets a default port mode (access)
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching**
- Configures Storm Control on all 10-Gigabit Ethernet network interfaces
- Provides basic Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP) configuration

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

- **802.3ad**—Specify an aggregated Ethernet bundle for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.
- **autonegotiation**—Enable or disable autonegotiation of flow control, link mode, and speed for interfaces.
- **link-mode**—Specify **full-duplex**, **half-duplex**, or **automatic** for Gigabit Ethernet interfaces.
- **loopback**—Enable or disable a loopback interface for both Gigabit Ethernet and 10-Gigabit Ethernet interfaces.

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

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

This topic describes:

- [Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches on page 60](#)
- [Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches on page 61](#)
- [Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches on page 61](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches on page 63](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches on page 63](#)
- [Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches on page 64](#)

Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches

If you are connecting a switch to other switches and to routers on the LAN, you need to assign the interface to a logical port and you need to configure the logical port as a trunk port.

To configure a Gigabit Ethernet or 10-Gigabit interface for trunk port mode on the Enhanced Layer 2 software (ELS):

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

See Also • [Monitoring Interface Status and Traffic on page 65](#)

Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches

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

The following default configurations are available on Gigabit Ethernet interfaces:

- You cannot set the speed on these interfaces.

On QFX5100-48S and QFX5100-96S devices using 1-Gigabit Ethernet SFP interfaces, the speed is set to 1 Gbps by default and cannot be configured to operate in a different speed.

- On QFX5100 devices, the interface naming for Gigabit Ethernet interfaces changes automatically to `xe-0/0/0`, `ge-0/0/0`, or `et-0/0/0` when the appropriate SFP is inserted.
- Gigabit Ethernet interfaces operate in full-duplex mode.
- Autonegotiation is supported by default. Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M and 1G. By default, autonegotiation is disabled on 10-Gigabit fiber ports.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



NOTE: Do not use the **set interfaces *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

Issue the **show interfaces *name* extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches

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

The following default configurations are available on Gigabit Ethernet interfaces:

- Gigabit Ethernet interfaces operate in full-duplex mode.
- Gigabit Ethernet interfaces must be configured as `xe-fpc/pic/port`, and not `ge-fpc/pic/port`.
- Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M, 1G, and 10G. By default, autonegotiation is enabled on 10-Gigabit Ethernet Copper interfaces.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



NOTE: Do not use the **set interfaces *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

You can reenabling autonegotiation by issuing the **set interfaces *name* ether-options auto-negotiate** command.

Issue the **show interfaces *name* extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches

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

- All the 10-Gigabit Ethernet interfaces are set to **auto-negotiation**.
- Flow control for 10-Gigabit Ethernet interfaces is set to **enabled** by default. You can disable flow control by specifying the **no-flow-control** option.
- The speed cannot be configured.

On QFX5100-48S, QFX5100-96S, and QFX5100-24Q devices using 10-Gigabit Ethernet SFP interfaces, the speed is set to 10 Gbps by default and cannot be configured to operate in a different speed.

- On QFX5100 devices, the interface naming for Gigabit Ethernet interfaces changes automatically to xe-0/0/0, ge-0/0/0, or et-0/0/0 when the appropriate SFP is inserted.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M and 1G. By default, autonegotiation is disabled on 10-Gigabit fiber ports.

If for some reason you have disabled autonegotiation, you can enable it by issuing the **set interfaces *name* ether-options auto-negotiate** command.

To disable autonegotiation, issue the **delete interfaces *name* ether-options auto-negotiate** command.



NOTE: Do not use the **set interfaces *name* ether-options no-auto-negotiate** command to remove the autonegotiation configuration.

Issue the **show interfaces *name* extensive** command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches

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

- All the 10-Gigabit Ethernet interfaces are set to **auto-negotiation**.
- Flow control for 10-Gigabit Ethernet interfaces is set to **enabled** by default. You can disable flow control by specifying the **no-flow-control** option.
- 10-Gigabit Ethernet interfaces operate in full-duplex mode by default.
- Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M, 1G, and 10G.



NOTE: In Junos OS Release 14.1X53-D35 on QFX5100-48T-6Q devices using 10-Gigabit Ethernet Copper interfaces, autonegotiation is disabled by default on the copper ports, and the interfaces operate at a speed of 100M. You can, however, enable auto-negotiation by issuing the `set interface name ether-options auto-negotiation` command on the interface for which you want to change the interface speed. With autonegotiation enabled, the interface auto-detects the speed in which to operate.

If for some reason you have disabled autonegotiation, you can enable it by issuing the `set interfaces name ether-options auto-negotiate` command.

Issue the `show interfaces name extensive` command to see if autonegotiation is enabled or disabled and the negotiated speed of the interface.

Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches

To specify an IP address for the logical unit:

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

Release History Table

| Release | Description |
|-------------|---|
| 14.1X53-D35 | In Junos OS Release 14.1X53-D35 on QFX5100-48T-6Q devices using 10-Gigabit Ethernet Copper interfaces, autonegotiation is disabled by default on the copper ports, and the interfaces operate at a speed of 100M. |

Related Documentation

- [Monitoring Interface Status and Traffic on page 65](#)
- [show interfaces xe on page 314](#)
- [show interfaces ge on page 294](#)
- [speed on page 213](#)

Configuring Ethernet Loopback Capability

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

```
loopback;
```

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

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

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

```
no-loopback;
```

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

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

Related Documentation

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)

Monitoring Interface Status and Traffic

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

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

Meaning For details about output from the CLI commands, see **show interfaces xe**.

Troubleshooting Network Interfaces

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

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

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

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

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

PART 3

Layer 3 Logical Interfaces

- [Understanding Layer 3 Logical Interfaces on page 69](#)

CHAPTER 3

Understanding Layer 3 Logical Interfaces

- [Understanding Layer 3 Logical Interfaces on page 69](#)
- [Configuring a Layer 3 Logical Interface on page 70](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 70](#)

Understanding Layer 3 Logical Interfaces

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



NOTE: You can also use Layer 3 logical interfaces to provide alternative gateway addresses for smart DHCP relay. The logical tunnel (lt) and virtual loopback tunnel (vt) interfaces are not supported in logical interfaces.

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

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

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

Configuring a Layer 3 Logical Interface

Devices use Layer 3 logical interfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. Layer 3 logical interfaces route traffic between subnets.

To configure Layer 3 logical interfaces, enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See *Configuring VLANs on Switches*.

To configure Layer 3 logical interfaces:

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-id-number
```

Verifying That Layer 3 Logical Interfaces Are Working

Purpose After configuring Layer 3 logical interfaces, verify that they are set up properly and transmitting data.

- Action**
1. To determine if you have successfully created the logical interfaces and the links are up:

```
[edit interfaces]
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 | 10.0.0.1/8 | |
| ge-0/0/0.1 | up | up | inet | 10.0.0.2/8 | |
| ge-0/0/0.2 | up | up | inet | 10.0.0.3/8 | |
| ge-0/0/0.3 | up | up | inet | 10.0.0.4/8 | |
| ge-0/0/0.4 | up | up | inet | 10.0.0.5/8 | |
| 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 if packets were transmitted correctly on the logical interface VLANs:

```
user@switch> ping ip-address
```

```
PING 10.1.1.1 (1.1.1.1): 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 logical interfaces have been created and the links are up.

PART 4

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

- [Understanding LAGs and LACP on page 75](#)

CHAPTER 4

Understanding LAGs and LACP

- Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80
- Configuring Link Aggregation on page 81
- Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85
- Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface on page 90
- Verifying the Status of a LAG Interface on page 90
- Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 91
- Troubleshooting an Aggregated Ethernet Interface on page 92

Understanding Aggregated Ethernet Interfaces and LACP for Switches

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.



NOTE: On QFX5100, EX4600, QFX10002 standalone switches, and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.



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

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



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



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

- [Link Aggregation Group on page 76](#)
- [Link Aggregation Control Protocol \(LACP\) on page 78](#)

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 20 on page 76](#) lists the EX Series switches and the maximum number of interfaces per LAG and the maximum number of LAGs they support.

Table 20: Maximum Interfaces per LAG and Maximum LAGs per Switch

| Switch | Maximum Interfaces per LAG | Maximum LAGs |
|-----------------------------------|----------------------------|--------------|
| EX2200 | 8 | 32 |
| EX2300 | 8 | 128 |
| EX3200 | 8 | 32 |
| EX3300 and EX3300 Virtual Chassis | 8 | 32 |
| EX3400 | 16 | 128 |
| EX4200 and EX4200 Virtual Chassis | 8 | 111 |
| EX4300 and EX4300 Virtual Chassis | 16 | 128 |

Table 20: Maximum Interfaces per LAG and Maximum LAGs per Switch (continued)

| Switch | Maximum Interfaces per LAG | Maximum LAGs |
|--|----------------------------|--------------|
| EX4500, EX4500 Virtual Chassis, EX4550, and EX4550 Virtual Chassis | 8 | 111 |
| EX4600 | 32 | 128 |
| EX6200 | 8 | 111 |
| EX8200 | 12 | 255 |
| EX8200 Virtual Chassis | 12 | 239 |
| EX9200 | 64 | 150 |

To create a LAG:

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

Keep in mind these hardware and software guidelines:

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



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

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and up to 448 LAGs are supported on QFX3500, QFX3600, EX4600, and OCX Series switches, and up to 1,000 LAGs are supported on QFX5100, QFX5200, QFX5110, QFX10002, QFX10008, and QFX10016 switches.



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

- Up to 64 Ethernet interfaces can be grouped to form a LAG, and in Junos Fusion, up to 1,000 LAGs are supported on QFX10002 switches acting as aggregation devices.
- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed and be in full-duplex mode.



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

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

Link Aggregation Control Protocol (LACP)

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

LACP was designed to achieve the following:

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

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 this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

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.

**Related
Documentation**

- [Verifying the Status of a LAG Interface on page 90](#)

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 for Switches” on page 85](#).

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

Before you configure LACP for EX Series, be sure you have:

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

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. Enable the LACP mode:

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

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

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



NOTE: LACP decides active and back up state of links. When configuring LACP, state of the backup link should not be configured manually as down. The following command is not supported if LACP is configured: `set interfaces ae0 aggregated-ether-options link-protection backup-state down`

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

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

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

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



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)*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85*
- *Configuring Aggregated Ethernet Interfaces (J-Web Procedure)*
- *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*

Configuring Link Aggregation

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



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



NOTE: On QFX5100, QFX5200, EX4600, QFX10002, and QFX10008 standalone switches and on QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated Ethernet bundle. Load balancing will not work if you configure link speeds that are not supported. (Platform support depends on the Junos OS release in your installation.)

1. [Creating an Aggregated Ethernet Interface on page 82](#)
2. [Configuring the VLAN Name and VLAN ID Number on page 83](#)
3. [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 83](#)

Creating an Aggregated Ethernet Interface

To create an aggregated Ethernet interface:

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

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

For example, to specify 5:

```
[edit chassis]
user@switch# set aggregated-devices interfaces device-count 5
```

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



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

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

For example, to specify 5:

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

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

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

For example, to specify 10g:

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

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

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

Configuring the VLAN Name and VLAN ID Number



NOTE: VLANs are not supported on OCX Series switches.

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

For example, 100.



NOTE: When you add or remove a vlan from a LAG interface, the interface goes down and comes back (flaps). The flapping happens when a low speed SFP is plugged into a relatively high speed port. To avoid flapping, configure the port speed to match the speed of the SFP.

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 for Switches” on page 85](#).

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

Before you configure LACP for EX Series, be sure you have:

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

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. Enable the LACP mode:

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

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

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



NOTE: LACP decides active and back up state of links. When configuring LACP, state of the backup link should not be configured manually as down. The following command is not supported if LACP is configured: `set interfaces ae0 aggregated-ether-options link-protection backup-state down`

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

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

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

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



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

- See Also**
- *Configuring Aggregated Ethernet Links (CLI Procedure)*
 - [Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85](#)
 - *Configuring Aggregated Ethernet Interfaces (J-Web Procedure)*
 - *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*

Related Documentation

- [Understanding Interface Naming Conventions on page 14](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch](#)
- [Verifying the Status of a LAG Interface on page 90](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 91](#)
- [show lacp statistics interfaces \(View\) on page 339](#)

Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches

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). For EX Series, see *Configuring Aggregated Ethernet Links (CLI Procedure)*.
- Configured LACP for the interface. For Ex Series, see “[Configuring Aggregated Ethernet LACP \(CLI Procedure\)](#)” on page 80.

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 87](#)
- [Configuring LACP Link Protection for a Single Link at the Aggregated Interface Level on page 87](#)
- [Configuring Subgroup Bundles to Provide LACP Link Protection to Multiple Links in an Aggregated Ethernet Interface on page 88](#)

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
```



NOTE: The LACP decides active and back up state of links. When configuring LACP, the state of the backup link should not be configured manually as down. The following command is not supported if LACP is configured: `set interfaces ae0 aggregated-ether-options link-protection backup-state down`

Related Documentation

- [lACP \(Aggregated Ethernet\) on page 248](#)

Configuring Periodic Rebalancing of Subscribers in an Aggregated Ethernet Interface

If subscribers are frequently logging in and logging out of your network, you can configure the system to periodically rebalance the links based on a specific time and interval.

To configure periodic rebalancing:

1. Access the aggregated Ethernet interface for which you want to configure periodic rebalancing.

```
edit
user@host# edit interfaces aenumber aggregated-ether-options
```

2. Configure the rebalancing parameters for the interface, including the time and the interval between rebalancing actions.

```
[edit interfaces aenumber aggregated-ether-options]
user@host# rebalance-periodic time hour:minute <interval hours>
```

Related Documentation

- [Verifying the Distribution of Demux Subscribers in an Aggregated Ethernet Interface](#)
- [Configuring the Distribution Type for Demux Subscribers on Aggregated Ethernet Interfaces](#)
- [Distribution of Demux Subscribers in an Aggregated Ethernet Interface](#)

Verifying the Status of a LAG Interface

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

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

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

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|-------|-------|--------|
| ae0 | up | up | | | |

```
ae0.0          up      up      inet    10.10.10.2/8
```

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

Related Documentation

- [Configuring Link Aggregation on page 81](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 91](#)
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*
- [show lacp statistics interfaces \(View\) on page 339](#)

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 91](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 92](#)

Verifying the LACP Setup

Purpose Verify that the LACP has been set up correctly.

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

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

```
Aggregated interface: ae0
LACP state:
xe-0/1/0      Actor No Yes No No No Yes Fast Active
xe-0/1/0      PartnerNo Yes No No No Yes Fast Passive
LACP protocol:
xe-0/1/0      Receive State Transmit State Mux State
                Defaulted Fast periodic Detached
```

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

Verifying That LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged between interfaces.

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

```
show lacp statistics interfaces ae0
```

```
Aggregated interface: ae0
```

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

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

Related Documentation

- [Verifying the Status of a LAG Interface on page 90](#)
- [show lacp statistics interfaces \(View\) on page 339](#)

Troubleshooting an Aggregated Ethernet Interface

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

Solution Check the following:

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



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

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

Related Documentation

- [Verifying the Status of a LAG Interface on page 90](#)
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*

PART 5

Load Balancing

- [Understanding Load Balancing on page 95](#)

CHAPTER 5

Understanding Load Balancing

- [Configuring Load Balancing Based on MAC Addresses on page 95](#)

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include the **multiservice** statement at the **[edit forwarding-options hash-key]** or **[edit chassis fpc slot number pic PIC number hash-key]** hierarchy level:

```
multiservice {
  source-mac;
  destination-mac;
  payload {
    ip {
      layer3-only;
      layer-3 (source-ip-only | destination-ip-only);
      layer-4;
      inner-vlan-id;
      outer-vlan-id;
    }
  }
}
```

To include the destination-address MAC information in the hash key, include the **destination-mac** option. To include the source-address MAC information in the hash key, include the **source-mac** option.



NOTE: Any packets that have the same source and destination address will be sent over the same path.



NOTE: You can configure per-packet load balancing to optimize EVPN traffic flows across multiple paths.



NOTE: Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.

Related Documentation

- *multiservice*

PART 6

Local Link Bias

- [Understanding Local Link Bias on page 99](#)

CHAPTER 6

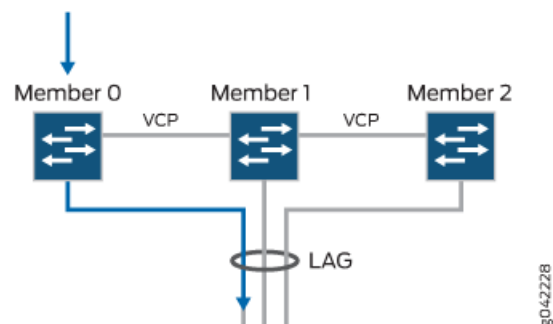
Understanding Local Link Bias

- [Understanding Local Link Bias on page 99](#)
- [Configuring Local Link Bias on page 101](#)

Understanding Local Link Bias


Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF using a different member link in the LAG bundle. The 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 99](#).

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 100](#).

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. |

Configuring Local Link Bias

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 **aex** 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
```


PART 7

Redundant Trunk Groups

- [Understanding Redundant Trunk Groups on page 105](#)

CHAPTER 7

Understanding Redundant Trunk Groups

- [Understanding Redundant Trunk Links \(Legacy RTG Configuration\) on page 106](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support on page 108](#)

Understanding Redundant Trunk Links (Legacy RTG Configuration)

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for network recovery when a trunk port on a switch goes down. In that case, traffic is routed to another trunk port, keeping network convergence time to a minimum.



NOTE: For information on redundant trunk link configurations that include Q-in-Q support and use LAGs with link protection, see *Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection*.

To configure a redundant trunk link, create a redundant trunk group. The redundant trunk group is configured on the access switch and contains two links: a primary or active link, and a secondary link. If the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal spanning-tree protocol convergence.

Data traffic is forwarded only on the active link. Data traffic on the secondary link is dropped and shown as dropped packets when you issue the operational mode command **show interfaces *interface-name* extensive**.

While data traffic is blocked on the secondary link, Layer 2 control traffic is still permitted. For example, an LLDP session can be run between two switches on the secondary link.

Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You must disable RSTP on an interface if a redundant trunk group is configured on that interface. For example, in [Figure 3 on page 107](#), in addition to disabling RSTP on the Switch 3 interfaces, you must also disable RSTP on the Switch 1 and Switch 2 interfaces connected to Switch 3. Spanning-tree protocols can, however, continue operating on other interfaces on those switches—for example on the link between Switch 1 and Switch 2.

[Figure 3 on page 107](#) shows three switches in a basic topology for redundant trunk links. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2). Link 1 and Link 2 are in a redundant trunk group called group1. Link 1 is designated as the primary link. Traffic flows between Switch 3 in the access layer and Switch 1 in the distribution layer through Link 1. While Link 1 is active, Link 2 blocks traffic.

Figure 3: Redundant Trunk Group, Link 1 Active





Figure 4 on page 107 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 4: Redundant Trunk Group, Link 2 Active



When Link 1 between Switch 1 and Switch 3 goes down, Link 2 takes over as the active link. Traffic between the access layer and the distribution layer is then automatically switched to Link 2 between Switch 3 and Switch 2.

Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support



NOTE: This example uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style.. For ELS details, see *Using the Enhanced Layer 2 Software CLI*.

You can manage network convergence by configuring both a primary link and a secondary link on a switch; this is called a redundant trunk group (RTG). If the primary link in a redundant trunk group fails, it passes its known MAC address locations to the secondary link, which automatically takes over after one minute.

This example describes how to create a redundant trunk group with a primary and a secondary link:

- [Requirements on page 108](#)
- [Overview and Topology on page 109](#)
- [Disabling RSTP on Switches 1 and 2 on page 111](#)
- [Configuring Redundant Trunk Links on Switch 3 on page 111](#)
- [Verification on page 112](#)

Requirements

This example uses the following hardware and software components:

- Two EX Series or QFX Series distribution switches
- One EX Series or QFX Series access switch
- The appropriate software release for your platform:
 - For EX Series switches: Junos OS Release 13.2X50-D10 or later
 - For the QFX Series: Junos OS Release 13.2X50-D15 or later

Before you configure the redundant trunk links network on the access and distribution switches, be sure you have:

- Configured interfaces ge-0/0/9 and ge-0/0/10 on the access switch, Switch 3, as trunk interfaces.
- Configured one trunk interface on each distribution switch, Switch 1 and Switch 2.
- Connected the three switches as shown in the topology for this example (see [Figure 5 on page 110](#)).

Overview and Topology

In a typical enterprise network composed of distribution and access layers, a redundant trunk link provides a simple solution for trunk interface network recovery. When a trunk interface fails, data traffic is routed to another trunk interface after one minute, thereby keeping network convergence time to a minimum.

This example shows the configuration of a redundant trunk group that includes one primary link (and its interface) and one unspecified link (and its interface) that serves as the secondary link.

A second type of redundant trunk group, not illustrated in the example, consists of two unspecified links (and their interfaces); in this case, neither of the links is primary. The software selects an active link by comparing the port numbers of the two links and activating the link with the higher port number. For example, if the two link interfaces use interfaces ge-0/1/0 and ge-0/1/1, the software activates ge-0/1/1. (In the interface names, the final number is the port number.)

The two links in a redundant trunk group generally operate the same way, whether they are configured as primary/unspecified or unspecified/unspecified. Data traffic initially passes through the active link but is blocked on the inactive link. While data traffic is blocked on the secondary link, note that Layer 2 control traffic is still permitted if the link is active. For example, an LLDP session can be run between two switches on the secondary link. If the active link either goes down or is disabled administratively, it broadcasts a list of its known MAC addresses for data traffic; the other link immediately picks up and adds the MAC addresses to its address table, becomes active, and begins forwarding traffic.

The one difference in operation between the two types of redundant trunk groups occurs when a primary link is active, goes down, is replaced by the secondary link, and then reactivates. When a primary link is re-enabled while the secondary link is active, the primary link waits 1 second (you can change the time interval by using the preempt cutover timer to accommodate your network) and then takes over as the active link. In other words, the primary link has priority and is always activated if it is available. This differs from the behavior of two unspecified links, both of which act as equals. Because the unspecified links are equal, the active link remains active until it either goes down or is disabled administratively; this is the only time that the other unspecified link learns the MAC addresses and immediately becomes active.

The example given here illustrates a primary/unspecified configuration for a redundant trunk group because that configuration gives you more control and is more commonly used.



NOTE: Rapid Spanning Tree Protocol (RSTP) is enabled by default on the switches to create a loop-free topology, but an interface is not allowed to be in both a redundant trunk group and in a spanning-tree protocol topology at the same time. You will need to disable RSTP on the two distribution switches in the example, Switch 1 and Switch 2. Spanning-tree protocols can, however, continue operating in other parts of the network—for example, between the distribution switches and also in links between distribution switches and the enterprise core.


Figure 5 on page 110 displays an example topology containing three switches. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk interfaces ge-0/0/9.0 (Link 1) and ge-0/0/10.0 (Link 2).

Table 21 on page 111 lists the components used in this redundant trunk group.

Because RSTP and RTGs cannot operate simultaneously on a switch, you disable RSTP on Switch 1 and Switch 2 in the first configuration task, and you disable RSTP on Switch 3 in the second task.

The second configuration task creates a redundant trunk group called example 1 on Switch 3. The trunk interfaces ge-0/0/9.0 and ge-0/0/10.0 are the two links configured in the second configuration task. You configure the trunk interface ge-0/0/9.0 as the primary link. You configure the trunk interface ge-0/0/10.0 as an unspecified link, which becomes the secondary link by default.

Figure 5: Topology for Configuring the Redundant Trunk Links



g020016

Table 21: Components of the Redundant Trunk Link Topology

| Property | Settings |
|-----------------------|--|
| Switch hardware | <ul style="list-style-type: none"> Switch 1–1 EX Series or QFX Series distribution switch Switch 2–1 EX Series or QFX Series distribution switch Switch 3–1 EX Series or QFX Series access switch |
| Trunk interfaces | On Switch 3 (access switch): ge-0/0/9.0 and ge-0/0/10.0 |
| Redundant trunk group | rtg0 |

Disabling RSTP on Switches 1 and 2

To disable RSTP on Switch 1 and Switch 2, perform this task on each switch:

CLI Quick Configuration To quickly disable RSTP on Switch 1 and Switch 2, copy the following command and paste it into each switch terminal window:

```
[edit]
set protocols rstp disable
```

Step-by-Step Procedure To disable RSTP on Switch 1 and Switch 2:

1. Disable RSTP on Switch 1 and Switch 2:

```
[edit]
user@switch# set protocols rstp disable
```

Results Check the results of the configuration:

```
[edit]
user@switch# show
protocols {
  rstp {
    disable;
  }
}
```

Configuring Redundant Trunk Links on Switch 3

To configure redundant trunk links on Switch 3, perform this task:

CLI Quick Configuration To quickly configure the redundant trunk group rtg0 on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp disable
set switch-options redundant-trunk-group group rtg0 interface ge-0/0/9.0 primary
set switch-options redundant-trunk-group group rtg0 interface ge-0/0/10.0
set switch-options redundant-trunk-group group rtg0 preempt-cutover-timer 60
```

Step-by-Step Procedure

Configure the redundant trunk group rtg0 on Switch 3.

1. Turn off RSTP:

```
[edit]
user@switch# set protocols rstp disable
```

2. Name the redundant trunk group rtg0 while configuring trunk interface ge-0/0/9.0 as the primary link and ge-0/0/10 as an unspecified link to serve as the secondary link:

```
[edit switch-options]
user@switch# set redundant-trunk-group group rtg0 interface ge-0/0/9.0 primary
user@switch# set redundant-trunk-group group rtg0 interface ge-0/0/10.0
```

3. (Optional) Change the time interval (from the default of 1 second) that a re-enabled primary link waits to take over for an active secondary link:

```
[edit switch-options]
user@switch# set switch-options redundant-trunk-group group rtg0 preempt-cutover-timer
60
```

Results Check the results of the configuration:

```
[edit]
user@switch# show
switch-options
  redundant-trunk-group {
    group rtg0 {
      preempt-cutover-timer 60;
      interface ge-0/0/9.0 {
        primary;
      }
      interface ge-0/0/10.0;
    }
  }
protocols {
  rstp {
    disable;
  }
}
```

Verification

To confirm that the configuration is set up correctly, perform this task:

- [Verifying That a Redundant Trunk Group Was Created on page 113](#)

Verifying That a Redundant Trunk Group Was Created

Purpose Verify that the redundant trunk group rtg0 has been created on Switch 1 and that trunk interfaces are members of the redundant trunk group.

Action List all redundant trunk groups configured on the switch:

```
user@switch> show redundant-trunk-group
```

| Group name | Interface | State | Time of last flap | Flap count |
|------------|-------------|--------|-------------------|------------|
| rtg0 | ge-0/0/9.0 | Up/Pri | Never | 0 |
| | ge-0/0/10.0 | Up | Never | 0 |

Meaning The `show redundant-trunk-group` command lists all redundant trunk groups configured on the switch as well as the interface names and their current states (up or down for an unspecified link, and up or down and primary for a primary link). For this configuration example, the output shows that the redundant trunk group rtg0 is configured on the switch. The **Up** beside the interfaces indicates that both link cables are physically connected. The **Pri** beside trunk interface ge-0/0/9.0 indicates that it is configured as the primary link.

PART 8

Configuration Statements and Operational Commands

- [Ethernet OAM Link Fault Management Configuration Statements on page 117](#)
- [Interfaces Configuration Statements on page 153](#)
- [LAGs and LACP Configuration Statements on page 233](#)
- [Redundant Trunk Groups Configuration Statements on page 253](#)
- [Ethernet OAM Link Fault Management Operational Command on page 259](#)
- [Interfaces Operational Commands on page 265](#)
- [LAGs and LACP Operational Commands on page 333](#)
- [Redundant Trunk Group Operational Command on page 341](#)

CHAPTER 8

Ethernet OAM Link Fault Management Configuration Statements

- [action \(OAM LFM\) on page 118](#)
- [action-profile on page 119](#)
- [allow-remote-loopback on page 120](#)
- [ethernet \(Protocols OAM\) on page 121](#)
- [event-thresholds on page 127](#)
- [event \(LFM\) on page 128](#)
- [frame-error on page 129](#)
- [frame-period on page 130](#)
- [frame-period-summary on page 131](#)
- [oam on page 132](#)
- [interface \(OAM Link-Fault Management\) on page 135](#)
- [link-adjacency-loss on page 136](#)
- [link-discovery on page 137](#)
- [link-down on page 137](#)
- [link-event-rate on page 138](#)
- [link-fault-management on page 139](#)
- [negotiation-options on page 140](#)
- [no-allow-link-events on page 140](#)
- [pdu-interval on page 141](#)
- [pdu-threshold on page 142](#)
- [remote-loopback on page 143](#)
- [symbol-period on page 144](#)
- [syslog \(OAM Action\) on page 145](#)
- [traceoptions \(Individual Interfaces\) on page 146](#)

action (OAM LFM)

| | |
|---------------------------------|--|
| Syntax | <pre>action { syslog; link-down; }</pre> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Define the action or actions to be taken when the OAM link fault management (LFM) fault event occurs.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management on page 4 |

action-profile

| | |
|---|--|
| List of Syntax | <p>Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches on page 119</p> <p>Syntax: EX Series Switches and NFX Series Devices on page 119</p> |
| <p>Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches</p> | <pre> action-profile <i>profile-name</i> { action { link-down; send-critical-event; syslog; } event { link-adjacency-loss; link-event-rate { frame-error <i>count</i>; frame-period <i>count</i>; frame-period-summary <i>count</i>; symbol-period <i>count</i>; } protocol-down; } } </pre> |
| <p>Syntax: EX Series Switches and NFX Series Devices</p> | <pre> action-profile <i>profile-name</i>; action { syslog; link-down; } event { link-adjacency-loss; link-event-rate { frame-error <i>count</i>; frame-period <i>count</i>; frame-period-summary <i>count</i>; symbol-period <i>count</i>; } } } </pre> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management] |
| Release Information | <p>Statement introduced in Junos OS Release 8.5 for T, M, MX and ACX Series Routers, SRX Series Firewalls, and EX Series Switches, .</p> <p>Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.</p> |
| Description | Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name. |

The remaining statements are explained separately. See [CLI Explorer](#).

Options *profile-name*—Name of the action profile.

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

Related Documentation

- [Configuring an OAM Action Profile](#)
- [Configuring Ethernet OAM Link Fault Management on page 4](#)

allow-remote-loopback

Syntax allow-remote-loopback;

Hierarchy Level [edit protocols [oam ethernet link-fault-management interface](#) *interface-name*]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Advertise that the interface is capable of getting into loopback mode. Enable remote loopback in Ethernet OAM link fault management (LFM) on all Ethernet interfaces or the specified interface on the EX Series switch.



WARNING: If you disable this statement on a peer interface, LFM loopback enable and disable commands will not work. Before disabling this configuration, please make sure the remote-loopback interface is disabled.

Required Privilege Level

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

Related Documentation

- [Example: Configuring Ethernet OAM Link Fault Management](#)
- [Configuring Ethernet OAM Link Fault Management on page 4](#)

ethernet (Protocols OAM)

List of Syntax [Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches on page 121](#)
[Syntax: EX Series Switches and NFX Series Devices on page 124](#)

Syntax: MX, T, ACX Series Routers, SRX Firewalls, M320 and EX Series Switches

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
    }
  }
  performance-monitoring {
    delegate-server-processing;
    hardware-assisted-timestamping;
    hardware-assisted-keepalives;
    sla-iterator-profiles {
      profile-name {
        avg-fd-twoway-threshold;
        avg-ifdv-twoway-threshold;
        avg-flr-forward-threshold;
        avg-flr-backward-threshold;
        disable;
        calculation-weight {
          delay delay-weight;
          delay-variation delay-variation-weight;
        }
        cycle-time milliseconds;
        iteration-period connections;
        measurement-type (loss | statistical-frame-loss | two-way-delay);
      }
    }
  }
}
linktrace {
  age (30m | 10m | 1m | 30s | 10s);
  path-database-size path-database-size;
}
maintenance-domain domain-name {
  level number;
  name-format (character-string | none | dns | mac+2octet);
  maintenance-association ma-name {
    short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    protect-maintenance-association protect-ma-name;
    remote-maintenance-association remote-ma-name;
    continuity-check {
      convey-loss-threshold;
      hold-interval minutes;
      interface-status-tlv;
      interval (10m | 10s | 1m | 1s | 100ms);
      loss-threshold number;
      port-status-tlv;
    }
  }
}
```

Copyright © 2019, Juniper Networks, Inc.

```
    }
    negotiation-options {
        allow-remote-loopback;
        no-allow-link-events;
    }
}
lmi {
    status-counter count;
    polling-verification-timer value;
    interface name {
        uni-id uni-name;
        status-counter number;
        polling-verification-timer value;
        evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
        evc evc-name {
            default-evc;
            vlan-list vlan-id-list;
        }
    }
}
}
```

**Syntax: EX Series
Switches and NFX
Series Devices**

```

ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      action {
        interface-down;
      }
      default-actions {
        interface-down;
      }
      event {
        adjacency-loss;
      }
    }
  }
  esp-traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable |
      no-world-readable>;
    flag (all | error | esp | interface | krt | lib | normal | task | timer);
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    level number;
    mip-half-function (none | default | explicit);
    name-format (character-string | none | dns | mac+2oct);
    maintenance-association ma-name {
      continuity-check {
        hold-interval minutes;
        interface-status-tlv;
        interval (10m | 10s | 1m | 1s | 100ms);
        loss-threshold number;
        port-status-tlv;
      }
      mep mep-id {
        auto-discovery;
        direction down;
        interface interface-name;
        priority
        remote-mep mep-id {
          action-profile profile-name;
          sla-iterator-profile profile-name {
            data-tlv-size size;
            iteration-count count-value;
            priority priority-value;
          }
        }
      }
    }
    short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
  }
}
performance-monitoring {
  sla-iterator-profiles {
    profile-name {
      calculation-weight {

```

125

```
}
```

Hierarchy Level [edit protocols oam]

Release Information Statement introduced in Junos OS Release 8.2 for MX, T, ACX Series routers, SRX firewalls, M320 and EX Series switches.
Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.

connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.

Description Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately. See [CLI Explorer](#).

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

Related Documentation

- *Enabling IEEE 802.3ah OAM Support*
- *Example: Configuring Ethernet OAM Link Fault Management*

event-thresholds

Syntax

```
event-thresholds {  
  frame-error count;  
  frame-period count;  
  frame-period-summary count;  
  symbol-period count;  
}
```

Hierarchy Level [edit protocols [oam ethernet link-fault-management interface](#) *interface-name*]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Configure threshold limit values for link events in periodic OAM PDUs.

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.

Related Documentation

- [Configuring Ethernet OAM Link Fault Management on page 4](#)

event (LFM)

List of Syntax [Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches on page 128](#)
[Syntax: EX Series Switches and NFX Series Devices on page 128](#)

Syntax: MX, M, T, ACX Series Routers, SRX Firewalls and EX Series Switches

```
event {
  link-adjacency-loss;
  link-event-rate {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  protocol-down;
}
```

Syntax: EX Series Switches and NFX Series Devices

```
event {
  link-adjacency-loss;
  link-event-rate {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
}
```

Hierarchy Level [edit protocols oam [ethernet](#) link-fault-management [action-profile](#)]

Release Information Statement introduced in Junos OS Release 8.5 for MX, M, T, ACX Series routers, SRX Series firewalls and EX Series switches.
 Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX devices.

Description Configure link events in an action profile for Ethernet OAM link fault management (LFM).
 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.
 routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • [Monitoring Protocol Status](#)
 • [Configuring Ethernet OAM Link Fault Management on page 4](#)

frame-error

| | |
|---------------------------------|--|
| Syntax | <code>frame-error <i>count</i>;</code> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management event link-event-rate], [edit protocols oam ethernet link-fault-management interface <i>interface-name</i> event-thresholds] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Configure the threshold value for sending frame error events or taking the action specified in the action profile.</p> <p>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.</p> |
| Options | <p><i>count</i>—Threshold count in seconds for frame error events.</p> <p>Range: 1 through 100 seconds</p> <p>Default: 1 second</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management on page 4 |

frame-period

| | |
|---------------------------------|--|
| Syntax | <code>frame-period <i>count</i>;</code> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management event link-event-rate], [edit protocols oam ethernet link-fault-management interface <i>interface-name</i> event-thresholds] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Configure the number of frame errors within the last N frames that has exceeded a threshold.</p> <p>Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.</p> |
| Options | <p><i>count</i>—Threshold count in seconds for frame error events.</p> <p>Range: 1 through 100 seconds</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management on page 4 |

frame-period-summary

| | |
|---------------------------------|---|
| Syntax | <code>frame-period-summary count;</code> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management event link-event-rate], [edit protocols oam ethernet link-fault-management interface interface-name event-thresholds] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Configure the threshold value for sending frame period summary error events or taking the action specified in the action profile.</p> <p>An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period.</p> |
| Options | <p><i>count</i>—Threshold count in seconds for frame period summary error events.</p> <p>Range: 1 through 100 seconds</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management on page 4 |

oam

```

Syntax  oam {
        ethernet {
            connectivity-fault-management {
                action-profile profile-name {
                    action {
                        interface-down;
                    }
                    default-actions {
                        interface-down;
                    }
                    event {
                        adjacency-loss;
                    }
                }
            }
            linktrace {
                age (30m | 10m | 1m | 30s | 10s);
                path-database-size path-database-size;
            }
            maintenance-domain domain-name {
                level number;
                mip-half-function (none | default | explicit);
                name-format (character-string | none | dns | mac+2oct);
                maintenance-association ma-name {
                    continuity-check {
                        hold-interval minutes;
                        interface-status-tlv;
                        interval (10m | 10s | 1m | 1s | 100ms);
                        loss-threshold number;
                        port-status-tlv;
                    }
                    mep mep-id {
                        auto-discovery;
                        direction down;
                        interface interface-name;
                        remote-mep mep-id {
                            action-profile profile-name;
                        }
                    }
                }
            }
            performance-monitoring {
                sla-iterator-profiles {
                    profile-name {
                        calculation-weight {
                            delay delay-value;
                            delay-variation delay-variation-value;
                        }
                        cycle-time cycle-time-value;
                        iteration-period iteration-period-value;
                        measurement-type two-way-delay;
                        passive;
                    }
                }
            }
        }
    }

```

```

    }
  }
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
}
}
}

```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.
connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.

Description Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.

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 Ethernet OAM Link Fault Management*
 - *Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches*
 - [Configuring Ethernet OAM Link Fault Management on page 4](#)
 - *Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)*

interface (OAM Link-Fault Management)

| | |
|---|--|
| List of Syntax | <p>Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches on page 135</p> <p>Syntax: EX Series Switches and NFX Series Devices on page 135</p> |
| <p>Syntax: T, M, MX and ACX Series Routers, SRX Series Firewalls and EX Series Switches</p> | <pre> interface <i>interface-name</i> { apply-action-profile <i>profile-name</i>; link-discovery (active passive); pdu-interval <i>interval</i>; pdu-threshold <i>threshold-value</i>; remote-loopback; event-thresholds { frame-error <i>count</i>; frame-period <i>count</i>; frame-period-summary <i>count</i>; symbol-period <i>count</i>; } negotiation-options { allow-remote-loopback; no-allow-link-events; } } </pre> |
| <p>Syntax: EX Series Switches and NFX Series Devices</p> | <pre> interface <i>interface-name</i> { link-discovery (active passive); pdu-interval <i>interval</i>; pdu-threshold <i>threshold-value</i>; remote-loopback; event-thresholds { frame-error <i>count</i>; frame-period <i>count</i>; frame-period-summary <i>count</i>; symbol-period <i>count</i>; } negotiation-options { allow-remote-loopback; no-allow-link-events; } } </pre> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management] |
| Release Information | <p>Statement introduced in Junos OS Release 8.2 for T, M, MX and ACX Series Routers, SRX Series firewalls and EX Series switches.</p> <p>Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.</p> |

| | |
|---------------------------------|---|
| Description | <p>Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.</p> <p>For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.3ah Operation, Administration, and Management (OAM) support.</p> |
| Options | <p>interface <i>interface-name</i>—Interface to be enabled for IEEE 802.3ah link fault management OAM support.</p> <p>The remaining statements are described separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• <i>Enabling IEEE 802.3ah OAM Support</i>• <i>Example: Configuring Ethernet OAM Link Fault Management</i>• Configuring Ethernet OAM Link Fault Management on page 4 |

link-adjacency-loss

| | |
|---------------------------------|---|
| Syntax | link-adjacency-loss; |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management action-profile event] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | Configure loss of adjacency event with the IEEE 802.3ah link fault management (LFM) peer. When included, the loss of adjacency event triggers the action specified under the action statement. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring Ethernet OAM Link Fault Management</i>• Configuring Ethernet OAM Link Fault Management on page 4 |

link-discovery

| | |
|---------------------------------|--|
| Syntax | link-discovery (active passive); |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | Specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on an interface. Link monitoring is done when the interface sends periodic OAM PDUs. |
| Options | <p><i>active</i>—In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.</p> <p><i>passive</i>—In passive mode, the peer initiates the discovery process.</p> <p>Once the discovery process is initiated, both sides participate in discovery.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management on page 4 |

link-down

| | |
|---------------------------------|--|
| Syntax | link-down; |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management action-profile action] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | Mark the interface as down for transit traffic. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management on page 4 |

link-event-rate

Syntax

```
link-event-rate {  
  frame-error count;  
  frame-period count;  
  frame-period-summary count;  
  symbol-period count;  
}
```

Hierarchy Level [edit protocols [oam ethernet link-fault-management action-profile event](#)]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Configure the number of link fault management (LFM) events per second.
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.

Related Documentation

- [Configuring Ethernet OAM Link Fault Management on page 4](#)

link-fault-management

```
Syntax link-fault-management {
    action-profile profile-name;
    action {
        syslog;
        link-down;
    }
    event {
        link-adjacency-loss;
        link-event-rate {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
        }
    }
    interface interface-name {
        link-discovery (active | passive);
        pdu-interval interval;
        pdu-threshold threshold-value;
        remote-loopback;
        event-thresholds {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
        }
        negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
        }
    }
}
```

Hierarchy Level [edit protocols [oam](#) [ethernet](#)]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

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.

Related Documentation

- *Example: Configuring Ethernet OAM Link Fault Management*
- [Configuring Ethernet OAM Link Fault Management on page 4](#)

negotiation-options

| | |
|---------------------------------|---|
| Syntax | <pre>negotiation-options { allow-remote-loopback; no-allow-link-events; }</pre> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Enable and disable IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) features for Ethernet interfaces.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management on page 4 |

no-allow-link-events

| | |
|---------------------------------|--|
| Syntax | <pre>no-allow-link-events;</pre> |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management interface <i>interface-name</i> negotiation-options] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | Disable the sending of link event TLVs. |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management on page 4 |

pdu-interval

| | |
|---------------------------------|--|
| Syntax | <code>pdu-interval <i>interval</i>;</code> |
| Hierarchy Level | <code>[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 8.2 for MX, M, T, ACX, Series routers, SRX Series firewalls, and EX Series Switches. Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. |
| Options | interval —Periodic OAM PDU sending interval. Range: For MX, M, T, ACX, Series routers, SRX Series firewalls and EX Series switches – 100 through 1000 milliseconds Default: For EX Series switches –1000 milliseconds |
| Required Privilege Level | <code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration. <code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring the OAM PDU Interval</i> • <i>Example: Configuring Ethernet OAM Link Fault Management</i> • Configuring Ethernet OAM Link Fault Management on page 4 |

pdu-threshold

| | |
|--------------------------|--|
| Syntax | <code>pdu-threshold <i>threshold-value</i>;</code> |
| Hierarchy Level | <code>[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]</code> |
| Release Information | <p>Statement introduced in Junos OS Release 8.2 for T, M, MX and ACX Series routers, SRX Series firewalls and EX Series switches.</p> <p>Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices.</p> |
| Description | <p>Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.</p> <p>For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.</p> |
| Options | <p><i>threshold-value</i>—The number of PDUs missed before declaring the peer lost.</p> <p>Range: 3 through 10 PDUs</p> <p>Default: 3 PDUs</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring the OAM PDU Threshold• Configuring Ethernet OAM Link Fault Management on page 4 |

remote-loopback

| | |
|---------------------------------|---|
| Syntax | remote-loopback; |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management interface <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | Set the data terminal equipment (DTE) in loopback mode. Remove the statement from the configuration to take the DTE out of loopback mode. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring Ethernet OAM Link Fault Management</i>• Configuring Ethernet OAM Link Fault Management on page 4 |

symbol-period

| | |
|---------------------------------|--|
| Syntax | <code>symbol-period <i>count</i>;</code> |
| Hierarchy Level | <code>[edit protocols oam ethernet link-fault-management action-profile; event link-event-rate]</code> , <code>[edit protocols oam ethernet link-fault-management interface <i>interface-name</i> event-thresholds]</code> |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Configure the threshold for sending symbol period events or taking the action specified in the action profile.</p> <p>Symbol code errors occur on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period. You cannot configure the default value to a different value.</p> |
| Options | <p><i>count</i>—Threshold count in seconds for symbol period events.</p> <p>Range: 1 through 100 seconds</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management on page 4 |

syslog (OAM Action)

| | |
|---------------------------------|--|
| Syntax | syslog; |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management action-profile action] |
| Release Information | Statement introduced in Junos OS Release 8.5 for T, M, MX and ACX Series routers, SRX Series firewalls and EX Series switches. Statement introduced in Junos OS Release 9.4 for EX Series switches and NFX Series devices. |
| Description | Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event. Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Specifying the Actions to Be Taken for Link-Fault Management Events</i> • Configuring Ethernet OAM Link Fault Management on page 4 |

traceoptions (Individual Interfaces)

| | |
|--|---|
| List of Syntax | Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) on page 146 Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) on page 146 Syntax (OAMLFM with EX Series, QFX Series, NFX Series) on page 146 Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) on page 146 |
| Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) | <pre>traceoptions { file <i>filename</i> <files <i>name</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; match; }</pre> |
| Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) | <pre>traceoptions { flag <i>flag</i>; }</pre> |
| Syntax (OAMLFM with EX Series, QFX Series, NFX Series) | <pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regex</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> ; no-remote-trace; }</pre> |
| Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) | <pre>traceoptions { file <<i>filename</i>> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; no-remote-trace; }</pre> |
| Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series) | [edit interfaces <i>interface-name</i>] |
| Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) | [edit interfaces] |

| | |
|----------------------------|--|
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in JUNOS Release 10.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>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 <i>/var/log/dcd</i>.</p> <p>On EX Series, QFX Series, and NFX Series platforms, configure tracing options the link fault management.</p> <p>On ACX Series, SRX Series, MX Series, M Series, and T Series platforms define tracing operations for the interface process (dcd).</p> |
| Default | If you do not include this statement, no interface-specific tracing operations are performed. |

Options [Table 22 on page 149](#) lists options for traceoption command for the following platforms:

Table 22: Options for traceoptions

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|----------------------|---|---|---|--|
| file 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 <code>/var/log/dcd</code> . By default, interface process tracing output is placed in the file. | | —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 <code>/var/log/dcd</code> . | —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 <code>/var/log/dcd</code> . By default, interface process tracing output is placed in the file <code>dcd</code> . |
| files number | —(Optional) Maximum number of trace files. When a trace file named <code>trace-file</code> 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. | | —(Optional) Maximum number of trace files. When a trace file named <code>trace-file</code> reaches its maximum size, it is renamed <code>trace-file.0</code> , then <code>trace-file.1</code> , and so on, until the maximum <code>xk</code> to specify KB, <code>xm</code> to specify MB, or <code>xg</code> to specify GB 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 <code>size</code> option. | —(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 <code>size</code> option. <i>Range: 2 through 1000</i> <i>Default: 3 files</i> |
| flag | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options. <ul style="list-style-type: none"> all—All interface tracing operations event—Interface events ipc—Interface interprocess | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options. <ul style="list-style-type: none"> all—All interface tracing operations event—Interface events ipc—Interface interprocess | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags: <ul style="list-style-type: none"> action-profile—Trace action profile invocation events. all—Trace all events. configuration—Trace configuration events. protocol—Trace | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags: <ul style="list-style-type: none"> all change-events—Log changes that produce configuration events config-states—Log the configuration |

Table 22: Options for traceoptions (continued)

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|------------------|---|---|--|--|
| | communication (IPC) messages <ul style="list-style-type: none"> • media—Interface media changes • q921—Trace ISDN Q.921 frames • q931—Trace ISDN Q.931 frames | communication (IPC) messages <ul style="list-style-type: none"> • media—Interface media changes • q921—Trace ISDN Q.921 frames • q931—Trace ISDN Q.931 frames | protocol processing events. <ul style="list-style-type: none"> • routing socket—Trace routing socket events. | state machine changes <ul style="list-style-type: none"> • kernel—Log configuration IPC messages to kernel • kernel-detail—Log details of configuration messages to kernel |
| match | —(Optional) Regular expression for lines to be traced. | | —(Optional) Refine the output to log only those lines that match the given regular expression. | |
| 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. | | —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). 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 files option. <p><i>Syntax:</i> xk to specify KB, xm to specify MB, or xg to specify GB</p> <p><i>Range:</i> 10 KB through 1 GB</p> <p><i>Default:</i> 128 KB</p> <p><i>Default:</i> If you do not include this option, tracing output is appended to an existing trace file.</p> | |

Table 22: Options for traceoptions (continued)

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|--------------------------|--|--|--|--|
| | | | | <p>—(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.</p> <p><i>Syntax:</i> xk to specify kilobytes, xm to specify megabytes, or xg to specify gigabytes</p> <p><i>Range:</i> 10 KB through the maximum file size supported on your router</p> <p><i>Default:</i> 1 MB</p> |
| no-world-readable | —(Optional) Prevent any user from reading the log file. | | —(Optional) Restrict file access to the user who created the file. | —(Optional) Disallow any user to read the log file. |
| world-readable | —(Optional) Allow any user to read the log file. | | —(Optional) Enable unrestricted file access. | —(Optional) Allow any user to read the log file. |
| disable | | | | |

Table 22: Options for traceoptions (continued)

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|--------------------|--|--|--|--|
| | | | | —(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 . |
| no-remote-trace | | | —(Optional) Disable the remote trace. | - |
| match <i>regex</i> | | | | —(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.
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

- Related Documentation**
- *Tracing Operations of an Individual Router Interface*
 - *Tracing Operations of an Individual Router or Switch Interface*
 - *Example: Configuring Ethernet OAM Link Fault Management*
 - [Configuring Ethernet OAM Link Fault Management on page 4](#)
 - *Tracing Operations of the Interface Process*

CHAPTER 9

Interfaces Configuration Statements

- [address on page 155](#)
- [alarm \(chassis\) on page 158](#)
- [auto-negotiation on page 159](#)
- [channel-speed on page 160](#)
- [configured-flow-control on page 161](#)
- [craft-lockout on page 162](#)
- [description \(Interfaces\) on page 163](#)
- [ethernet \(Alarm\) on page 164](#)
- [ethernet-switch-profile on page 165](#)
- [ethernet-switching on page 167](#)
- [ether-options on page 168](#)
- [eui-64 on page 169](#)
- [family on page 170](#)
- [fibre-channel \(Alarm\) on page 174](#)
- [filter on page 175](#)
- [flow-control on page 177](#)
- [fpc on page 178](#)
- [gratuitous-arp-reply on page 179](#)
- [hold-time \(Physical Interface\) on page 180](#)
- [irb \(Interfaces\) on page 182](#)
- [inet \(interfaces\) on page 185](#)
- [inet6 \(interfaces\) on page 186](#)
- [interface-mode on page 187](#)
- [interface-range on page 189](#)
- [interfaces \(QFX Series\) on page 191](#)
- [link-down on page 198](#)
- [link-mode on page 199](#)
- [link-speed on page 200](#)

- [loopback \(Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet\) on page 202](#)
- [mac on page 203](#)
- [management-ethernet \(Alarm\) on page 204](#)
- [member on page 205](#)
- [member-range on page 206](#)
- [mtu on page 207](#)
- [no-gratuitous-arp-request on page 208](#)
- [pic on page 209](#)
- [rx-buffers on page 211](#)
- [source on page 212](#)
- [speed \(Ethernet\) on page 213](#)
- [targeted-broadcast on page 219](#)
- [traceoptions \(Individual Interfaces\) on page 220](#)
- [tx-buffers on page 227](#)
- [unit on page 229](#)
- [vlan-id on page 230](#)
- [vlan-tagging on page 231](#)

address

```

Syntax  address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        master-only;
        multipoint-destination address dlcid dlcid-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
                 length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        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.



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 57](#).

- 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.

Options *address*—Address of the interface.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.



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

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


Related Documentation

- *Configuring the Protocol Family*
- *family*
- *negotiate-address*
- *unnumbered-address (Ethernet)*

alarm (chassis)

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


auto-negotiation

| | |
|--------------------------|--|
| Syntax | (auto-negotiation no-auto-negotiation); |
| Hierarchy Level | [edit interfaces <i>interface-name</i> ether-options] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Explicitly enable or disable autonegotiation. Autonegotiation is enabled by default, and will autonegotiate the speed with the link partner. We recommend that you keep autonegotiation enabled for interfaces operating at 100M, 1G, and 10G. |
| | <p>.....</p> <div>  <p>NOTE: In Junos OS Release 14.1X53-D35 on QFX5100-48T-6Q devices using 10-Gigabit Ethernet Copper interfaces, autonegotiation is disabled by default on the copper ports, and the interfaces operate at a speed of 100M. You can, however, enable auto-negotiation by issuing the <code>set interface <i>name</i> ether-options auto-negotiation</code> command on the interface for which you want to change the interface speed. With autonegotiation enabled, the interface auto-detects the speed in which to operate.</p> </div> <p>.....</p> <ul style="list-style-type: none"> • auto-negotiation—Enable autonegotiation. • no-auto-negotiation—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure link mode and speed options. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • speed on page 213 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

channel-speed

| | |
|---------------------------------|---|
| Syntax | <code>channel-speed (10g 25g 50g; 100g disable-auto-speed-detection) ;</code> |
| Hierarchy Level | <code>[edit chassis fpc slot-number pic pic-number (port port-number port-range port-range-low port-range-high)]</code> |
| Release Information | 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 | —Enable the specified port on the Physical Interface Card (PIC) to perform in the specified channel speed. Additionally, you can disable auto-speed detection. |
| Default | 40g (40-Gigabit Ethernet). |
| Options | 10g —Set the channel speed to 10g (10-Gigabit Ethernet). 25g —Set the channel speed to 25g (25-Gigabit Ethernet). 50g —Set the channel speed to 50g (50-Gigabit Ethernet). 100g —Set the channel speed to 100g (100-Gigabit Ethernet). disable-auto-speed-detection —Disable auto-speed detection. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Channelizing Interfaces Overview</i>• <i>Channelizing Interfaces on QFX5200-32C Switches</i> |

configured-flow-control

| | |
|---------------------------------|---|
| Syntax | <pre>configured-flow-control { rx-buffers (on off); tx-buffers (on off); }</pre> |
| Hierarchy Level | [edit interfaces interface-name ether-options] |
| Release Information | Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | <p>Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the rx-buffers and the tx-buffers values when you configure asymmetric flow control.</p> <p>Use the flow-control and no-flow-control statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> |
| Default | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces. |
| Options | The remaining statements are explained separately. See CLI Explorer . |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> congestion-notification-profile flow-control on page 177 |

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

craft-lockout

Syntax

```
craft-lockout {
  alarm {
    interface-type {
      link-down (red | yellow | ignore);
    }
  }
  container-devices {
    device-count number;
  }
  fpc slot {
    pic pic-number {
      fibre-channel {
        port-range {
          port-range-low port-range-high;
        }
      }
    }
  }
  routing-engine
    on-disk-failure {
      disk-failure-action (halt | reboot);
    }
  }
}
```

Hierarchy Level [edit chassis -interconnect-device]

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

Description Disable the physical operation of the craft interface front panel.

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

Related Documentation

- *Configuring the Junos OS to Disable the Physical Operation of the Craft Interface*

description (Interfaces)

| | |
|---------------------------------|---|
| Syntax | <code>description text;</code> |
| Hierarchy Level | <code>[edit interfaces interface-name],</code> <code>[edit interfaces interface-name unit logical-unit-number],</code> <code>[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]</code> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>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.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p> |
| Options | text —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Interface Description</i> • <i>Adding a Logical Unit Description to the Configuration</i> • <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i> • <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches</i> • <i>Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support</i> • <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59</i> • <i>Using DHCP Relay Agent Option 82 Information</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> • <i>Example: Connecting Access Switches with ELS Support to a Distribution Switch with ELS Support</i> |

ethernet (Alarm)

| | |
|---------------------------------|--|
| Syntax | <pre>ethernet { link-down (red yellow ignore); }</pre> |
| Hierarchy Level | <pre>[edit chassis alarm], [edit chassis interconnect-device name alarm], [edit chassis node-group name alarm]</pre> |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure alarms for an Ethernet interface. |
| Options | The remaining statements are explained separately. Search for a statement in CLI Explorer or click a linked statement in the Syntax section for details. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |

ethernet-switch-profile

Syntax

```
ethernet-switch-profile {
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [values];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    premium {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
  }
  storm-control storm-control-profile;
  tag-protocol-id tpid;
}
mac-learn-enable;
```

Hierarchy Level [edit interfaces *interface-name* gigether-options],
 [edit interfaces *interface-name* aggregated-ether-options],
 [edit interfaces *interface-name* **aggregated-ether-options**],
 [edit interfaces *interface-name* ether-options]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.
 Statement introduced in Junos OS Release 13.2 for the QFX Series.
 Statement introduced in Junos OS Release 13.2X50-D15 for the EX Series switches.

Description



NOTE: On QFX Series standalone switches, the **ethernet-policer-profile** CLI hierarchy and the **mac-learn-enable** statement are supported only on the Enhanced Layer 2 Switching CLI.

For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, the built-in Gigabit Ethernet port on the M7i router); 100-Gigabit Ethernet Type 5 PIC with CFP; and Gigabit Ethernet, 10-Gigabit Ethernet, 40-Gigabit Ethernet, and aggregated Ethernet interfaces on EX Series switches, configure VLAN tag and MAC address accounting and filtering properties.

The remaining statements are explained separately. See [CLI Explorer](#).



NOTE: When you gather interfaces into a bridge domain, the `no-mac-learn-enable` statement at the [edit interfaces *interface-name* *gather-options* ethernet-switch-profile] hierarchy level is not supported. You must use the `no-mac-learning` statement at the [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*] hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

| | |
|---------------------------------|--|
| Default | If the <code>ethernet-switch-profile</code> statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring Gigabit Ethernet Policers</i>• <i>Configuring MAC Address Filtering</i>• <i>Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview</i>• <i>Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support</i> |

ethernet-switching

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

Hierarchy Level [edit [interfaces](#) *ge-chassis/slot/port unit logical-unit-number*] family

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure Ethernet switching protocol family information for the logical interface.
The remaining statements are explained separately. See [CLI Explorer](#).

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

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

Related Documentation

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
- [JUNOS Software Network Interfaces Configuration Guide](#)

ether-options

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

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

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

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
autostate-exclude option introduced in Junos OS Release 14.1x53-D40 for QFX5100 switches only.

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



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

The remaining statements are explained separately. See [CLI Explorer](#).

Default Enabled.

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

- Related Documentation**
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
 - *Junos OS Network Interfaces Library for Routing Devices*

eui-64

| | |
|---------------------------------|--|
| Syntax | eui-64; |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i>] |
| 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 57 |

family

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

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

```

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

```

```

        priority-hold-time seconds;
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-inet6-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-group group-name;
        active-interface interface-name;
    }
}
(dad-disable | no-dad-disable);
filter {
    group filter-group-number;
    input filter-name;
    input-list [ filter-names ];
    output filter-name;
    output-list [ filter-names ];
}
mtu bytes;
nd6-stale-time time;
no-neighbor-learn;
no-redirects;
policer {
    input policer-name;
    output policer-name;
}
rpf-check {
    fail-filter filter-name;
    mode {
        loose;
    }
}
mpls {
    filter {
        group filter-group-number;
        input filter-name;
        input-list [ filter-names ];
        output filter-name;
        output-list [ filter-names ];
    }
    mtu bytes;
}
}
}

```

Hierarchy Level

[edit [interfaces interface-name unit logical-unit-number](#)],
 [edit [interfaces interface-range interface-name unit logical-unit-number family](#)]

Release Information

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

Description Configure protocol family information for the logical interface on the QFX Series and OCX Series product.

Default



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

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

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

Options

Interface types on the switch are:

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



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

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

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

Required Privilege Level

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


Related Documentation

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
- [Configuring Link Aggregation on page 81](#)
- [Configuring IRB Interfaces on Switches](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

fibre-channel (Alarm)


| | |
|---------------------------------|--|
| Syntax | <pre>fibre-channel { link-down (red yellow ignore); }</pre> |
| Hierarchy Level | <pre>[edit chassis alarm], [edit chassis interconnect-device name alarm], [edit chassis node-group name alarm]</pre> |
| Release Information | Statement introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Configure alarms for a Fibre Channel interface. |
| Options | The remaining statement is explained separately.— |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |

filter

| | |
|--------------------------|---|
| Syntax | <pre>filter { group <i>filter-group-number</i>; input <i>filter-name</i>; input-list [<i>filter-names</i>]; output <i>filter-name</i>; output-list [<i>filter-names</i>]; }</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. |
| Description | <p> NOTE: On EX Series switches, the <code>group</code>, <code>input-list</code>, <code>output-filter</code> statements are not supported under the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet6], and [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchies.</p> <p>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.</p> |
| Options | <p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. The default filter group number is 0. Range: 0 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p> |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Applying a Filter to an Interface</i> |

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

- Related Documentation**
- [configured-flow-control on page 161](#)
 - [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
 - *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*
 - *Junos OS Network Interfaces Library for Routing Devices*

fpc

Syntax

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

Hierarchy Level [edit chassis]

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

Description Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.

For generic routing encapsulation (GRE) tunneling, use the **tunnel-port** statement to specify the port that you want to convert to a GRE tunnel port.

Options *slot*—Number of the FPC slot. For QFX3500, QFX3600, QFX5200, and OCX Series devices, the slot number is always 0.

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

- *show chassis fpc*

gratuitous-arp-reply

| | |
|---------------------------------|---|
| Syntax | (gratuitous-arp-reply no-gratuitous-arp-reply); |
| Hierarchy Level | [edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-range</i> <i>interface-range-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Enable processing of ARP updates received via gratuitous ARP reply messages. |
| Default | Updating of the ARP cache is disabled on all Ethernet interfaces. |
| Options | gratuitous-arp-reply —Update the ARP cache. no-gratuitous-arp-reply —Do not update the ARP cache. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |

hold-time (Physical Interface)

| | |
|---------------------|---|
| Syntax | hold-time up <i>milliseconds</i> down <i>milliseconds</i> ; |
| 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 Metro 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. |
| | <p> NOTE:</p> <ul style="list-style-type: none"> 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. |
| | <p> 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.</p> |
| 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 interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related • *advertise-interval*
Documentation • *interfaces (EX Series switches)*
 • *Physical Interface Damping Overview*
 • *Damping Shorter Physical Interface Transitions*
 • *Damping Longer Physical Interface Transitions*

irb (Interfaces)

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

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

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



```

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

```

```

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

```

| | |
|---------------------------------|--|
| Hierarchy Level | [edit interfaces <i>interface-name</i> |
| Release Information | Statement introduced in Junos OS Release 12.3R2 for EX Series switches. irb option introduced in Junos OS Release 13.2 for the QFX Series. |
| Description | Configure the properties of a specific integrated bridging and routing (IRB) 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. |



inet (interfaces)

| | |
|---------------------------------|---|
| Syntax | <pre>inet { address <i>address</i> { primary; filter input <i>filter-name</i>; filter output <i>filter-name</i>; targeted-broadcast; } }</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family], [edit interfaces interface-range <i>interface-name</i> unit <i>logical-unit-number</i> family] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure the primary IP address for the logical interface. |
| Default | You must configure a logical interface to be able to use the physical device. |
| 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 | <ul style="list-style-type: none"> • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59 |

inet6 (interfaces)

| | |
|---------------------------------|--|
| Syntax | <pre>inet6 { address address { eui-64 preferred primary; filter input <i>filter-name</i>; filter output <i>filter-name</i>; } }</pre> |
| Hierarchy Level | [edit interfaces interface-name unit logical-unit-number family], [edit interfaces interface-range interface-name unit logical-unit-number family] |
| Release Information | 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 | Configure the primary IP address for the logical interface. |
| Default | You must configure a logical interface to be able to use the physical device. |
| 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 | <ul style="list-style-type: none"> • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches |

interface-mode

| | |
|---------------------|---|
| Syntax | <code>interface-mode (access trunk <inter-switch-link>);</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> 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 | <p> NOTE: This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see <i>port-mode</i>. For ELS details, see <i>Using the Enhanced Layer 2 Software CLI</i>.</p> <p>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.</p> <p> 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 <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> |
| Options | <p>access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the vlan-id statement.</p> <p>trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the vlan-id or vlan-id-list statement.</p> |

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 interface—To view this statement in the configuration.
Level 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 with ELS Support to a Distribution Switch with ELS Support*
 - *Tunnel Services Overview*
 - *Tunnel Interface Configuration on MX Series Routers Overview*

interface-range

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

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

Hierarchy Level [edit [interfaces](#)]

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

Description



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



NOTE: The interface range definition is supported only for Gigabit Ethernet, 10-Gigabit Ethernet, and Fibre Channel interfaces. Interface ranges are not supported on channelized interfaces.

Group interfaces that share a common configuration profile.

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



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

The remaining statements are explained separately. 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**
- [Understanding Interface Ranges for Switches on page 27](#)
 - [Interfaces Overview for Switches on page 9](#)
 - [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
 - *Junos OS Network Interfaces Library for Routing Devices*

interfaces (QFX Series)

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

```

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

```

```

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

```

```

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

```

```

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

```

```

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

```

```

virtual-address [ addresses ];
}
xe-0/0/port {
  disable;
  description text;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
        (primary | backup);
      }
    }
    configured-flow-control {
      rx-buffers (on | off);
      tx-buffers (on | off);
    }
    flexible-vlan-tagging;
    (flow-control | no-flow-control);
    loopback;
    no-loopback;
  }
  (gratuitous-arp-reply | no-gratuitous-arp-reply)
  hold-time milliseconds down milliseconds;
  mac
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no traps);
  unit logical-unit-number {
    disable;
    description text;
    family {
      ethernet-switching {
        filter input filter-name;
        filter output filter-name;
        native-vlan-id vlan-id;
        port-mode mode;
        reflective-relay;
        vlan {
          members [ (all | names | vlan-ids) ];
        }
      }
      fibre-channel {
        port-mode (f-port | np-port);
      }
      inet {
        address address {
          primary;
        }
        filter input filter-name;
        filter output filter-name;
        primary;
        targeted-broadcast;
      }
    }
    (traps | no traps);
  }
}

```

```

    vlan-id vlan-id-number;
  }
  vlan-tagging;
}

```

Hierarchy Level [\[edit\]](#)

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

Description Configure the interfaces on the QFX Series and OCX Series.

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

Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.

Options **aex**—Configure an aggregated Ethernet interface.

xe-0/0/*port***/**—Configure a 10-Gigabit Ethernet interface.

ge-0/0/*port***/**—Configure a Gigabit Ethernet interface.

fc-0/0/*port***/**—Configure a Fibre Channel interface.

meX/—Configure a management interface.

mc-ae—Configure a multichassis aggregated Ethernet (MC-AE) 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

- [Interfaces Overview for Switches on page 9](#)
- [Understanding Interface Ranges for Switches on page 27](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
- [Configuring Link Aggregation on page 81](#)
- [Configuring a Layer 3 Logical Interface on page 70](#)

link-down

| | |
|---------------------------------|--|
| Syntax | <code>link-down (red yellow ignore);</code> |
| Hierarchy Level | <code>[edit chassis alarm ethernet (Alarm)],</code> <code>[edit chassis alarm fibre-channel],</code> <code>[edit chassis interconnect-device <i>name</i> alarm ethernet (Alarm)],</code> <code>[edit chassis node-group <i>name</i> alarm fibre-channel]</code> |
| Release Information | Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify either red, yellow, or ignore to display when the link is down. |
| Options | <p>red—Indicates that one or more hardware components have failed or exceeded temperature thresholds, or an alarm condition configured on an interface has triggered a critical warning.</p> <p>yellow—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.</p> <p>ignore—Suppresses or ignores the alarm.</p> |
| Required Privilege Level | <code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration. |

link-mode

| | |
|--------------------------|---|
| Syntax | <code>link-mode mode;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> ether-options] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Set the device's link-connection characteristic. |
| | <div>  <p>NOTE: link-mode configuration is not supported on QFX10000 line of switches.</p> </div> |
| Default | The full-duplex mode is enabled. |
| Options | <p>mode —Link characteristic:</p> <ul style="list-style-type: none"> • full-duplex—Connection is full duplex. • half-duplex—Connection is half duplex. • automatic—Link mode is negotiated. <p>If no-auto-negotiation is specified in the ether-options option, you can select only full-duplex or half-duplex. If auto-negotiation is specified in the ether-options option, you can select any mode.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

link-speed

| | |
|---|--|
| Syntax | <code>link-speed <i>speed</i>;</code> |
| Hierarchy Level (QFX, NFX, EX Series, QFabric System, OCX1100, EX4600) | [edit interfaces aex aggregated-ether-options] |
| Hierarchy Level (EX Series) | [edit interfaces aex aggregated-ether-options], [edit interfaces interface-range <i>name</i> aggregated-ether-options], [edit interfaces interface-range <i>name</i> aggregated-sonet-options] |
| Release Information | Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | For aggregated Ethernet interfaces only, set the required link speed. |
| Options | <p><i>speed</i>—For aggregated Ethernet links, you can specify <i>speed</i> 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).</p> <p>Aggregated Ethernet links on the M120 router can have one of the following speeds:</p> <ul style="list-style-type: none"> • 100m—Links are 100 Mbps. • 10g—Links are 10 Gbps. • 1g—Links are 1 Gbps. • oc192—Links are OC192 or STM64c. <p>Aggregated Ethernet links on EX Series switches can be configured to operate at one of the following speeds:</p> <ul style="list-style-type: none"> • 10m—Links are 10 Mbps. • 100m—Links are 100 Mbps. • 1g—Links are 1 Gbps. • 10g—Links are 10 Gbps. <p><i>speed</i>—For aggregated Ethernet links, you can specify the speed in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).</p> <p>On QFX5100 and EX4600 standalone switches and on a QFX5100 Virtual Chassis and EX4600 Virtual Chassis, you can configure a mixed rate of link speeds for the aggregated</p> |

Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.

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.




NOTE: OCX Series switches only support 10g and 40g interfaces. Mixed rate aggregated Ethernet interfaces are not support on the OCX Series.

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

Related Documentation

- [Configuring Link Aggregation on page 81](#)
- [Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)
- [Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch](#)

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

| | |
|----------------------------|--|
| Syntax | (loopback no-loopback); |
| Hierarchy Level | <p>[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> gigether-options], [edit interfaces interface-range <i>name</i> ether-options]</p> <p>For QFX Series and EX Series:</p> <p>[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> ether-options],</p> <p>For SRX Series Devices and vSRX:</p> <p>[edit interfaces <i>interface-name</i> redundant-ether-options]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4 for MX Series.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Statement modified in Junos OS Release 9.2 for the SRX Series.</p> |
| Description | For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode. |
| | <p> NOTE:</p> <ul style="list-style-type: none"> 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 <i>ge-fpc/pic/port</i> gigether-options] hierarchy level, an NDP address cannot be configured at the [edit interfaces <i>ge-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet6 address] hierarchy level. |
| Default | By default, loopback is disabled. |

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*
- *Understanding Interfaces*

mac

Syntax `mac mac-address;`

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the MAC address of the interface.

Use this statement at the [edit interfaces ... ps0] hierarchy level to configure the MAC address for a pseudowire logical device that is used for subscriber interfaces over point-to-point MPLS pseudowires.

Options *mac-address*—MAC 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.

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

Related Documentation

- *Configuring the MAC Address on the Management Ethernet Interface*
- *Configuring a Pseudowire Subscriber Logical Interface Device*

management-ethernet (Alarm)

Syntax `management-ethernet {
 link-down (red | yellow | ignore);
 }`

Hierarchy Level `[edit chassis alarm],`
 `[edit chassis interconnect-device name alarm],`
 `[edit chassis node-group name alarm]`

Release Information 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 Configure alarms for a management Ethernet interface.



NOTE: If you configure a yellow alarm on the Interconnect device, it will be handled as a red alarm.

Options The remaining statement is explained separately.—

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

member

| | |
|---------------------------------|---|
| Syntax | <code>member <i>interface-name</i>;</code> |
| Hierarchy Level | [edit interfaces interface-range <i>interface-range-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the name of the member interface belonging to an interface range on the QFX Series switch. |
| Options | <i>interface-name</i> —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 | <ul style="list-style-type: none"> • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches • Interfaces Overview for Switches on page 9 • Junos OS Network Interfaces Library for Routing Devices |

member-range

| | |
|---------------------------------|--|
| Syntax | <code>member-range <i>starting-interface-name ending-interface-name</i>;</code> |
| Hierarchy Level | [edit interfaces interface-range <i>interface-range-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the names of the first and last members of a sequence of interfaces belonging to an interface range. |
| Options | <i>starting interface-name ending interface-name</i> —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 | <ul style="list-style-type: none">• Understanding Interface Ranges for Switches on page 27• Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59• Configuring Gigabit and 10-Gigabit Ethernet Interfaces for OCX Series Switches• Interfaces Overview for Switches on page 9• Interfaces Overview for Switches on page 9• Junos OS Network Interfaces Library for Routing Devices |

mtu

| | |
|---------------------|--|
| Syntax | <code>mtu bytes;</code> |
| Hierarchy Level | <code>[edit interfaces interface-name],</code> <code>[edit interfaces interface-range interface-name]</code> |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | <p>Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On QFX3500, QFX3600, QFX5100, and OCX Series switches, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight.</p> <p>On QFX5110 switches, setting MTU on the L3 interface does not take effect and packets with MTU greater than the default value are dropped.</p> <p>Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:</p> <ul style="list-style-type: none"> • For LAG interfaces—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (aex) automatically configures the jumbo MTU size on the member links. • For RVIs—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named vlan. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the vlan interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named vlan (the RVI). On a QFX5100 switch jumbo frames on the RVI are configured on the basis of the interface MTU. |



NOTE: RVIs are not supported on OCX Series switches.



CAUTION: Setting or deleting the jumbo MTU size on the RVI (the **vlan** interface) while the switch is transmitting packets might result in dropped packets.

| | |
|---------------------------------|---|
| Options | bytes —MTU size. Range: 64 through 9216 bytes Default: 1514 bytes |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

no-gratuitous-arp-request

| | |
|---------------------------------|---|
| Syntax | no-gratuitous-arp-request; |
| Hierarchy Level | [edit interfaces interface-name], [edit interfaces interface-range interface-name] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs). |
| Default | Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring IRB Interfaces on Switches</i> |

pic

| | |
|---|---|
| List of Syntax | Syntax (EX4600, OCX1100, QFX Series, QFabric system with ELS) on page 209 Syntax (EX4600, QFX Series, QFabric system) on page 209 |
| Syntax (EX4600, OCX1100, QFX Series, QFabric system with ELS) | <pre> pic <i>pic-number</i> { tunnel-port <i>port-number</i> tunnel-services; port <i>port-number</i> { channel-speed (<i>speed</i> disable-auto-speed-detection) ; } port-range <i>port-range-low</i> <i>port-range-high</i> { channel-speed (<i>speed</i> disable-auto-speed-detection) ; } } </pre> |
| Syntax (EX4600, QFX Series, QFabric system) | <pre> pic <i>pic-number</i> { fibre-channel { port-range { <i>port-range-low</i> <i>port-range-high</i>; } } fte { port <i>port-number</i>; (<i>port-range</i> <i>port-range-low</i> <i>port-range-high</i>); } xe { (<i>port</i> <i>port-number</i> <i>port-range</i> <i>port-range-low</i> <i>port-range-high</i>); } xle { (<i>port</i> <i>port-number</i> <i>port-range</i> <i>port-range-low</i> <i>port-range-high</i>); } } </pre> |
| Hierarchy Level (EX4600, OCX1100, QFX Series, QFabric system with ELS) | [edit chassis (EX Series) fpc <i>slot</i>] |
| Hierarchy Level (EX4600, QFX Series, QFabric system) | [edit chassis (QFX Series) <i>fpc</i> <i>slot</i>] [edit chassis (QFX Series) <i>node-group</i> <i>name</i> <i>node-device</i> <i>name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Options xe and xle introduced in Junos OS 12.2X50-D20 for the QFX Series. Option channel-speed introduced in Junos OS Release 13.2 for the QFX Series. |

- Description** (QFX3500, QFX3600, and QFX5100 standalone switches running Enhanced Layer 2 Software only)—Configure a specific port or a range of ports to operate as 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.
- Enable the specified port on the Physical Interface Card (PIC) to perform in the specified operating mode.
- Options** **pic *pic-number***—(QFX3500 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **1** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.
- (QFX3600 standalone switch only) Number of the physical interface card (PIC) on which you want to configure port types. Specify **0** to configure 10-Gigabit Ethernet or 40-Gigabit Ethernet type ports.
- port *physical-port-number***—Port number on which you want to configure the port type.
- port-range *port-range-low***—Lowest-numbered port in the range of ports.
- port-range *port-range-high***—Highest-numbered port in the range of ports.
- channel-speed (*speed* | **disable-auto-speed-detection**)** —Configure *10g* for 10-Gigabit Ethernet type ports, and configure **disable-auto-speed-detection** to disable auto-channelization.



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


- On a QFX3500 switch, specify **0** if the port type is **fiber-channel**, and **2** if the port type is **xle**.
- On a QFX3600 switch, specify **0** if the port type is **xe**, and **1** if the port type is **xle**.
- On a QFX5100 switch, specify **0** if the port type is **xe**, and **1** if the port type is **xle** and **fte**.

The remaining statements are explained separately. See [CLI Explorer](#).

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

- Related Documentation**
- *Channelizing Interfaces Overview*
 - *Configuring the QSFP+ Port Type on QFX3500 Standalone Switches*
 - *Configuring the Port Type on QFX3600 Standalone Switches*
 - *Configuring the QSFP+ Port Type on QFX5100 Devices*

rx-buffers

| | |
|----------------------------|---|
| Syntax | <code>rx-buffers (on off);</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> ether-options configured-flow-control] |
| Release Information | Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | <p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the rx-buffers statement with the tx-buffers statement to configure asymmetric Ethernet PAUSE on an interface. (Use the flow-control statement to enable symmetric PAUSE and the no-flow-control statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> |
| Default | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces. |
| Options | on off —Enable or disable an interface to generate and send Ethernet PAUSE messages. |

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

Related Documentation

- [flow-control on page 177](#)
- [tx-buffers on page 227](#)
- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*

source

| | |
|--|---|
| Syntax | <code>source <i>source-address</i>;</code> |
| Hierarchy Level (EX, NFX, OCX1100 and QFX Series) | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]</code> |
| Hierarchy Level (M-series and T-series) | <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel <i>address</i>]</code> |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the source address of the tunnel. |
| Default | If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel. |
| Options | <i>source-address</i> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • <i>Tunnel Services Overview</i> • <i>multicast-only</i> • <i>primary (Address on Interface)</i> • <i>Junos OS Services Interfaces Library for Routing Devices</i> |

speed (Ethernet)

| | |
|---|---|
| List of Syntax | Syntax (EX Series) on page 213 Syntax (EX2300 and EX4300) on page 213 Syntax (EX Series, ACX Series, MX Series) on page 213 Syntax (QFX Series, OCX1100, EX4600) on page 213 |
| Syntax (EX Series) | <code>speed (auto-negotiation <i>speed</i>) ;</code> |
| Syntax (EX2300 and EX4300) | <code>speed <i>speed</i>;</code> |
| Syntax (EX Series, ACX Series, MX Series) | <code>speed (10m 10g 100m 1g 2.5g 5g auto auto-10m-100m);</code> |
| Syntax (QFX Series, OCX1100, EX4600) | <code>speed (10g 1g 100m)</code> |
| Hierarchy Level (EX Series) | <code>[edit interfaces <i>interface-name</i> ether-options]</code> |
| Hierarchy Level (EX2300 and EX4300) | <code>[edit interfaces <i>interface-name</i>]</code> |
| Hierarchy Level (ACX Series, EX Series, MX Series) | <code>[edit interfaces <i>interface-name</i>],</code> <code>[edit interfaces ge-<i>pim</i>/0/0 switch-options switch-port <i>port-number</i>]</code> |
| Hierarchy Level (QFX Series, EX4600, OCX Series) | <code>[edit interfaces <i>interface-name</i>]</code> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.</p> <p>Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches.</p> <p>Speed option 2.5Gbps introduced in Junos OS Release 18.1R2 for EX2300 switch.</p> <p>Speed option 10Gbps and 5Gbps introduced in Junos OS Release 18.2R1 for EX4300 switch.</p> <p>Speed option 1-Gbps is introduced in Junos OS Release 19.1R1 on the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches.</p> |
| 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.

On 10-Gigabit Ethernet SFP interfaces, autonegotiation is enabled by default and auto-detects the speed to be either 1 Gbps or 10 Gbps. On QFX5100-48S, QFX5100-96S, and QFX5100-24Q devices using 10-Gigabit Ethernet SFP interfaces, the speed is set to 10 Gbps by default and cannot be configured to operate in a different speed. On QFX5100-48S and QFX5100-96S devices using 1-Gigabit Ethernet SFP interfaces, the speed is set to 1 Gbps by default and cannot be configured to operate in a different speed.



NOTE: In Junos OS Release 14.1X53-D35 on QFX5100-48T-6Q devices using 10-Gigabit Ethernet Copper interfaces, autonegotiation is disabled by default on the copper ports, and the interfaces operate at a speed of 100M. You can, however, enable auto-negotiation by issuing the `set interface name ether-options auto-negotiation` command on the interface for which you want to change the interface speed. With autonegotiation enabled, the interface auto-detects the speed in which to operate.



NOTE: Only 10 Gbps and 40 Gbps interfaces are supported on OCX Series switches.



NOTE: When displaying interface information with `show interfaces` commands, you might see speed values for 1 Gbps interfaces displayed as 1000mbps.

(For EX2300 only) Starting in Junos OS Release 18.1R2, the multi-rate speed is supported on EX2300-48MP and EX2300-24MP switches. The speed configuration statement is supported on both multi-rate gigabit ethernet interface (mge) and gigabit ethernet (ge) interface. The mge interface is a rate-selectable (multirate) Gigabit Ethernet interface that can support speeds of 10-Gbps, 5-Gbps, and 2.5-Gbps over CAT5e/CAT6/CAT6a cables. In the EX2300, the mge interface supports 100-Mbps, 1-Gbps, and 2.5-Gbps speeds, which can be configured by using the speed configuration statement. Note that 10Mbps speed is supported only on **ge** interfaces of EX2300 switch.

On EX2300-24MP and EX2300-48MP switches, if both Energy Efficient Ethernet (EEE) and 100-Mbps speed are configured on a rate-selectable (or multirate) Gigabit Ethernet (mge) port, the port operates only at 100-Mbps speed but EEE is not enabled on that port. EEE is supported only on mge interfaces that operate at 1-Gbps and 2.5-Gbps speeds.

(For EX4300-48MP only) Starting with Junos OS Release 19.1R1, the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module (EX-UM-4SFPP-MR) on EX4300-48MP switches supports 1-Gbps speed. You do not need to explicitly configure 1-Gbps speed on the uplink module as it automatically identifies the installed 1-gigabit SFP transceivers and creates the interface accordingly.



NOTE: On EX4300-48MP, the status LED of 1-Gigabit Ethernet uplink module port is solid green (instead of blinking green) because of a device limitation. However, there is no impact on device functionality.

Default (EX Series) If the **auto-negotiation** statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled, the auto-negotiation option is enabled by default.

Options 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:

Table 23: Options for speed

| Platforms | Speed Supported | Auto-negotiation |
|---|---|---|
| EX Series Switches | 100m —100 Mbps | auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the [edit interfaces interface-name ether-options] hierarchy level is enabled. |
| | 10m —10 Mbps | |
| | 1g —1 Gbps | |
| ACX, MX Series | 100m —100 Mbps | auto —Automatically negotiate the speed (10 Mbps, 100 Mbps, or 1 Gbps) based on the speed of the other end of the link. |
| | 10m —10 Mbps | |
| | 1g —1 Gbps | auto-10m-100m —Automatically negotiate the speed (10 Mbps or 100 Mbps) based on the speed of the other end of the link. |
| EX4600, QFX Series, QFabric, OCX100, QFX Series | 10g —10 Gbps | auto-negotiation—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the auto-negotiation statement at the [edit interfaces interface-name ether-options] hierarchy level is enabled. |
| | 1g —1 Gbps | |
| | 100m —100 Mbps | |
| EX2300 | 10m —10 Mbps (supported on EX series switches and only on ge interfaces of EX2300 switch) | speed —Specify the interface speed. If the auto-negotiation statement at the [edit interfaces interface-name ether-options] hierarchy level is disabled, you must specify a specific value. This value sets the speed that is used on the link. If the auto-negotiation statement is enabled, you might want to configure a specific speed value to advertise the desired speed to the remote end. The Multi-rate gigabit ethernet interface (MGE) on EX2300-24MP and EX2300-48MP switches flaps (becomes unavailable, and then available again) while performing timeout detection and recovery (TDR) test. |
| | 100m —100 Mbps | |
| | 1g —1 Gbps | |
| | 2.5g —2.5 Gbps (supported only on mge interfaces of E2300 switch) | |
| EX4300-48MP (EX-UM-4SFPP-MR) | 10m —10 Mbps (supported only on ge interfaces) | speed —Specify the interface speed. NOTE: On 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module, no explicit configuration is required as it automatically identifies the transceivers and creates the interface accordingly. |
| | 100m —100 Mbps (supported on ge and mge interfaces) | |
| | 1g —1 Gbps (supported on ge , mge interfaces, and 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches). The 1-Gbps speed is supported on the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module of EX4300-48MP switches from Junos OS Release 19.1R1 onwards. | |
| | 2.5g —2.5 Gbps (supported only on mge interfaces) | |
| | 5g —5 Gbps (supported only on mge interfaces) | |
| | 10g —10 Gbps (supported on mge interfaces) | |

Table 23: Options for speed (continued)

| |
|--|
| and 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module on EX4300-48MP switches) |
|--|

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

Release History Table

| Release | Description |
|---------|---|
| 18.2R1 | Starting in Junos OS Release 18.1R2, the multi-rate speed is supported on EX2300-48MP and EX2300-24MP switches. |

Related Documentation

- *Configuring the Interface Speed*
- *Configuring the Interface Speed on Ethernet Interfaces*
- *Configuring Gigabit Ethernet Autonegotiation*
- *Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support*
- [auto-negotiation on page 159](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59](#)
- *Junos OS Network Interfaces Library for Routing Devices*
- *Configuring Gigabit Ethernet Interfaces (CLI Procedure)*
- *Configuring Gigabit Ethernet Interfaces (J-Web Procedure)*
- [Junos OS Ethernet Interfaces Configuration Guide](#)

targeted-broadcast

| | |
|---|---|
| Syntax (EX Series, MX Series, ACX Series) | <pre>targeted-broadcast { forward-and-send-to-re; forward-only; }</pre> |
| Syntax (QFX Series, OCX1100, EX4600, NFX Series) | <pre>targeted-broadcast;</pre> |
| Hierarchy Level (EX Series, MX Series, ACX Series) | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]</pre> |
| Hierarchy Level (QFX Series, OCX1100, EX4600, NFX Series) | <pre>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit interfaces interface-range <i>interface-range-name</i> unit <i>logical-unit-number</i> family inet]</pre> |
| Release Information | <p>Statement introduced in Junos OS Release 9.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Specify the IP packets destined for a Layer 3 broadcast address to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p> |
| Default | When this statement is not included, broadcast packets are sent to the Routing Engine only. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Targeted Broadcast</i> • <i>Understanding Targeted Broadcast</i> |

traceoptions (Individual Interfaces)

| | |
|--|---|
| List of Syntax | Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) on page 220 Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) on page 220 Syntax (OAMLFM with EX Series, QFX Series, NFX Series) on page 220 Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) on page 220 |
| Syntax (Individual interfaces with PTX Series, EX Series, ACX Series) | <pre>traceoptions { file <i>filename</i> <files <i>name</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; match; }</pre> |
| Syntax (Individual interfaces with QFX Series, OCX1100, EX4600, NFX Series) | <pre>traceoptions { flag <i>flag</i>; }</pre> |
| Syntax (OAMLFM with EX Series, QFX Series, NFX Series) | <pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regex</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> ; no-remote-trace; }</pre> |
| Syntax (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) | <pre>traceoptions { file <<i>filename</i>> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <disable>; no-remote-trace; }</pre> |
| Hierarchy Level (Individual interfaces with PTX Series, EX Series, ACX Series, QFX Series, OCX1100, EX4600, NFX Series) | [edit interfaces <i>interface-name</i>] |
| Hierarchy Level (Interface process with ACX Series, SRX Series, MX Series, M Series, T Series) | [edit interfaces] |

| | |
|----------------------------|--|
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro Routers.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in JUNOS Release 10.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Define tracing operations for individual interfaces.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p> <p>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 <i>/var/log/dcd</i>.</p> <p>On EX Series, QFX Series, and NFX Series platforms, configure tracing options the link fault management.</p> <p>On ACX Series, SRX Series, MX Series, M Series, and T Series platforms define tracing operations for the interface process (dcd).</p> |
| Default | If you do not include this statement, no interface-specific tracing operations are performed. |

Options [Table 22 on page 149](#) lists options for traceoption command for the following platforms:

Table 24: Options for traceoptions

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|----------------------|---|---|---|--|
| file 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 <code>/var/log/dcd</code> . By default, interface process tracing output is placed in the file. | | —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 <code>/var/log/dcd</code> . | —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 <code>/var/log/dcd</code> . By default, interface process tracing output is placed in the file <code>dcd</code> . |
| files number | —(Optional) Maximum number of trace files. When a trace file named <code>trace-file</code> 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. | | —(Optional) Maximum number of trace files. When a trace file named <code>trace-file</code> reaches its maximum size, it is renamed <code>trace-file.0</code> , then <code>trace-file.1</code> , and so on, until the maximum <code>xk</code> to specify KB, <code>xm</code> to specify MB, or <code>xg</code> to specify GB 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 <code>size</code> option. | —(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 <code>size</code> option. <i>Range: 2 through 1000</i> <i>Default: 3 files</i> |
| flag | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options. <ul style="list-style-type: none"> all—All interface tracing operations event—Interface events ipc—Interface interprocess | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. The following are the interface-specific tracing options. <ul style="list-style-type: none"> all—All interface tracing operations event—Interface events ipc—Interface interprocess | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags: <ul style="list-style-type: none"> action-profile—Trace action profile invocation events. all—Trace all events. configuration—Trace configuration events. protocol—Trace | —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags: <ul style="list-style-type: none"> all change-events—Log changes that produce configuration events config-states—Log the configuration |

Table 24: Options for traceoptions (continued)

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|------------------|---|---|--|--|
| | communication (IPC) messages <ul style="list-style-type: none"> • media—Interface media changes • q921—Trace ISDN Q.921 frames • q931—Trace ISDN Q.931 frames | communication (IPC) messages <ul style="list-style-type: none"> • media—Interface media changes • q921—Trace ISDN Q.921 frames • q931—Trace ISDN Q.931 frames | protocol processing events. <ul style="list-style-type: none"> • routing socket—Trace routing socket events. | state machine changes <ul style="list-style-type: none"> • kernel—Log configuration IPC messages to kernel • kernel-detail—Log details of configuration messages to kernel |
| match | —(Optional) Regular expression for lines to be traced. | | —(Optional) Refine the output to log only those lines that match the given regular expression. | |
| 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. | | —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). 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 files option. <p><i>Syntax:</i> xk to specify KB, xm to specify MB, or xg to specify GB</p> <p><i>Range:</i> 10 KB through 1 GB</p> <p><i>Default:</i> 128 KB</p> <p><i>Default:</i> If you do not include this option, tracing output is appended to an existing trace file.</p> | |

Table 24: Options for traceoptions (continued)

| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|--------------------------|--|--|--|--|
| | | | | <p>—(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.</p> <p><i>Syntax:</i> xk to specify kilobytes, xm to specify megabytes, or xg to specify gigabytes</p> <p><i>Range:</i> 10 KB through the maximum file size supported on your router</p> <p><i>Default:</i> 1 MB</p> |
| no-world-readable | —(Optional) Prevent any user from reading the log file. | | —(Optional) Restrict file access to the user who created the file. | —(Optional) Disallow any user to read the log file. |
| world-readable | —(Optional) Allow any user to read the log file. | | —(Optional) Enable unrestricted file access. | —(Optional) Allow any user to read the log file. |
| disable | | | | |

Table 24: Options for traceoptions (continued)


| Option | Individual interfaces with PTX Series, ACX Series, EX Series | Individual interfaces with QFX Series, QFabric System, OCX1100, EX4600, NFX Series | Interface Process with OAMLFM with EX Series, QFX Series, NFX Series | Interface process with ACX Series, SRX Series, MX Series, M Series, T Series |
|--------------------|--|--|--|--|
| | | | | —(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 . |
| no-remote-trace | | | —(Optional) Disable the remote trace. | - |
| match <i>regex</i> | | | | —(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.
 routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

- Related Documentation**
- *Tracing Operations of an Individual Router Interface*
 - *Tracing Operations of an Individual Router or Switch Interface*
 - *Example: Configuring Ethernet OAM Link Fault Management*
 - [Configuring Ethernet OAM Link Fault Management on page 4](#)
 - *Tracing Operations of the Interface Process*

tx-buffers

| | |
|--------------------------|---|
| Syntax | tx-buffers (on off); |
| Hierarchy Level | [edit interfaces <i>interface-name</i> ether-options configured-flow-control] |
| Release Information | Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | <p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the tx-buffers statement with the rx-buffers statement to configure asymmetric Ethernet PAUSE on an interface. (Use the flow-control statement to enable symmetric PAUSE and the no-flow-control statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> <div>  <p>NOTE: Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> |
| Default | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces. |
| Options | on off —Enable or disable an interface to respond to an Ethernet PAUSE message. |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |

**Related
Documentation**

- [flow-control on page 177](#)
- [rx-buffers on page 211](#)
- *Configuring CoS Asymmetric Ethernet PAUSE Flow Control*
- *Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*

unit

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

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

Hierarchy Level [edit [interfaces \(QFX Series\) interface-name](#)],
[edit [interfaces \(QFX Series\) interface-range interface-range-name](#)]

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

Description



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

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


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

Options ***logical-unit-number***—Number of the logical unit.
Range: 0 through 16,384

The remaining statements are explained separately. 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 | <ul style="list-style-type: none"> • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59 • Configuring Link Aggregation on page 81 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

vlan-id

| | |
|---------------------------------|--|
| Syntax | <code>vlan-id <i>vlan-id-number</i>;</code> |
| Hierarchy Level | [edit interfaces (QFX Series) <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 9.2 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | For 10-Gigabit Ethernet and aggregated Ethernet interfaces only, bind an 802.1Q VLAN tag ID to a logical interface. Statement introduced in Junos OS Release 9.2 for EX Series switches. |
| | <div>  <p>NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.</p> </div> |
| Options | <i>vlan-id-number</i> —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 | <ul style="list-style-type: none"> • vlan-tagging on page 231 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces for EX4600 and QFX Series Switches on page 59 • Configuring a Layer 3 Logical Interface on page 70 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

vlan-tagging

| | |
|---|---|
| Syntax | vlan-tagging; |
| Syntax (QFX Series, NFX Series, and EX4600) | vlan-tagging; |
| Syntax (SRX Series Interfaces) | vlan-tagging native-vlan-id <i>vlan-id</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>] |
| QFX Series, NFX Series, and EX4600 Interfaces | [edit interfaces (QFX Series) <i>interface-name</i>] [edit interfaces (QFX Series) <i>interface-range</i> <i>interface-range-name</i>] |
| SRX Series Interfaces | [edit interfaces <i>interface</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Metro 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: For QFX Series configure VLAN identifier for untagged packets received on the physical interface of a trunk mode interface. Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags.

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.

Default VLAN tagging is disabled by default.

Options **native-vlan-id**— (SRX Series) Configures a VLAN identifier for untagged packets. Enter a number from 0 through 4094.



NOTE: The **native-vlan-id** can be configured only when either **flexible-vlan-tagging mode** or **interface-mode trunk** is configured.

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*
 - *Configuring a Layer 3 Subinterface (CLI Procedure)*
 - *Configuring Tagged Aggregated Ethernet Interfaces*
 - *Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch*
 - [vlan-id on page 230](#)
 - [Configuring a Layer 3 Logical Interface on page 70](#)
 - *Configuring VLAN Tagging*

CHAPTER 10

LAGs and LACP Configuration Statements

- [aggregated-devices on page 234](#)
- [aggregated-ether-options on page 236](#)
- [chassis \(QFX Series\) on page 239](#)
- [802.3ad on page 241](#)
- [device-count on page 243](#)
- [ethernet on page 244](#)
- [force-up on page 245](#)
- [lacp \(802.3ad\) on page 246](#)
- [lacp \(Aggregated Ethernet\) on page 248](#)
- [link-protection on page 251](#)
- [periodic on page 252](#)

aggregated-devices

List of Syntax [Syntax \(EX Series\) on page 234](#)
[Syntax \(QFX Series, EX4600, OCX1100, NFX Series\) on page 234](#)

Syntax (EX Series)

```
aggregated-devices {
  ethernet (Aggregated Devices) {
    device-count number;
    lacp
  }
}
```

Syntax (QFX Series, EX4600, OCX1100, NFX Series)

```
aggregated-devices {
  ethernet {
    device-count number;
  }
}
```

Hierarchy Level (EX Series, QFX Series) [\[edit chassis \(EX Series\)\]](#)

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
 Statement introduced in Junos OS Release 14.2R3

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

- [Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75](#)
- [Configuring Link Aggregation on page 81](#)
- *Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch*
- *Junos OS Network Interfaces Library for Routing Devices*
- *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)*

- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

aggregated-ether-options

List of Syntax [Syntax \(EX, MX Series\) on page 236](#)
 [Syntax \(NFX, QFX Series, EX4600, OCX1100, QFabric\) on page 236](#)

Syntax (EX, MX Series)

```

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
}

```

Syntax (NFX, QFX Series, EX4600, OCX1100, QFabric) The `fcoe-lag` and `mc-ae` statements are not supported on OCX Series switches.

```

aggregated-ether-options {
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  ethernet-switch-profile {
    tag-protocol-id;
    (fcoe-lag | no-fcoe-lag);
    (flow-control | no-flow-control);
  }
}

```

```

lacp mode {
    admin-key key;
    periodic interval;
    system-id mac-address;
    force-up;
}
(link-protection | no-link-protection);
link-speed speed;
local-bias;
local-minimum-links-threshold threshold-value;
(loopback | no-loopback);
mc-ae {
    chassis-id chassis-id;
    mc-ae-id mc-ae-id;
    mode (active-active);
    status-control (active | standby);
}
minimum-links number;
rebalance-periodic;
resilient-hash;
source-address-filter filter;
(source-filtering | no-source-filtering);
}

```

Hierarchy Level (EX Series, QFX Series)

[edit interfaces *aex*]

Release Information

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 12.3R2.
Statements **fcoe-lag** and **no-fcoe-lag** introduced in Junos OS Release 13.2X52-D10 for the QFX Series.
Statements **force-up**, **lacp**, and **resilient-hash** introduced in Junos OS Release 14.1X53-D10 for the QFX Series.
Statement **local-minimum-links-threshold** introduced in Junos OS Release 14.1X53-D40 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

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



NOTE:

- The `fcoe-lag` and `mc-ae` statements are not supported on OCX Series switches.
- The `force-up` statement is not supported on QFX10002 switches.
- The `resilient-hash` statement is not supported on QFX10002 switches.

The remaining statements are explained separately. See [CLI Explorer](#).

Default Options are not enabled.

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

Related Documentation

- [Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 80](#)
- *Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch*
- *Junos OS Network Interfaces Library for Routing Devices*
- *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 \(CLI Procedure\)](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 80](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85](#)
- [Configuring Q-in-Q Tunneling on EX Series Switches with ELS Support](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

chassis (QFX Series)

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

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Statement introduced in Junos OS Release 14.2R3

Description Configure chassis-specific properties for the switch.


The remaining statements are explained separately. See [CLI Explorer](#).

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

Related Documentation

- *Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion*
- [Configuring Link Aggregation on page 81](#)

802.3ad

| | |
|--|--|
| Syntax (EX Series) | <pre>802.3ad { aex; (backup primary); lacp { force-up; port-priority } }</pre> |
| Syntax (NFX, OCX, and QFX Series) | <pre>802.3ad ae{x; lacp { force-up; (primary backup); } port-priority; }</pre> |
| Hierarchy Level (EX Series) | [edit interfaces <i>interface-name</i> ether-options] |
| Hierarchy Level (NFX, OCX, and QFX Series) | [edit <i>interfaces</i> <i>interface-name</i> <i>ether-options</i>] |
| Release Information | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Specify the aggregated Ethernet logical interface number. |
| <div>  <p>NOTE:</p> <ul style="list-style-type: none"> The <i>port-priority</i> statement is not supported on QFabric systems. The force-up statement is not supported on QFX10002 switches. </div> | |
| Options | <ul style="list-style-type: none"> aex—Name of the LAG. Aggregated Ethernet logical interface number. backup—Designate the interface as the backup interface for link-protection mode. primary—Designate the interface as the primary interface for link-protection mode. <p>The remaining statements are described separately. See CLI Explorer.</p> |

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i>• <i>Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch</i>• <i>Example: Configuring Multicast Load Balancing for Use with Aggregated 10-Gigabit Ethernet Interfaces on EX8200 Switches</i>• <i>Configuring Aggregated Ethernet Links (CLI Procedure)</i>• Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80• Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85• Configuring Link Aggregation on page 81• Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80• Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75• Troubleshooting an Aggregated Ethernet Interface on page 92• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

device-count

| | |
|---|--|
| Syntax (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System) | <code>device-count <i>number</i>;</code> |
| Hierarchy Level (EX Series) | [edit chassis (EX Series) aggregated-devices ethernet (Aggregated Devices)] |
| Hierarchy Level (EX, NFX, QFX Series, EX4600, OCX1100, QFabric System) | [edit chassis aggregated-devices ethernet], [edit chassis node-group <i>name</i> aggregated-devices ethernet] |
| 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. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 14.2R3 |
| Description | Configure the number of aggregated Ethernet logical devices available to the switch. |
| Options | <p><i>number</i>—Maximum number of aggregated Ethernet logical interfaces on the switch.</p> <p>Range: 1 through 32 for EX2200, EX3200, and standalone EX3300 switches and for EX3300 Virtual Chassis</p> <p>Range: 1 through 64 for standalone EX4200, standalone EX4500, and EX6200 switches and for EX4200 and EX4500 Virtual Chassis</p> <p>Range: 1 through 239 for EX8200 Virtual Chassis</p> <p>Range: 1 through 255 for standalone EX8200 switches</p> <p>Range: 1 through 480 for standalone EX9200 switches</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion • Configuring Link Aggregation on page 81 • Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch • 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) • Junos OS Network Interfaces Configuration Guide |



ethernet

| | |
|---------------------------------|--|
| Syntax | <pre>ethernet { device-count <i>number</i>; }</pre> |
| Hierarchy Level | [edit chassis aggregated-devices], [edit chassis node-group aggregated-devices] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 14.2R3 |
| Description | Configure properties for aggregated Ethernet 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 | <ul style="list-style-type: none">• <i>Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion</i>• Configuring Link Aggregation on page 81• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

force-up

| | |
|---------------------------------|---|
| Syntax | <code>force-up;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> ether-options 802.3ad lacp] |
| Release Information | Statement introduced in Junos OS Release 10.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. Statement introduced in Junos OS Release 16.1 for EX9200 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 | <ul style="list-style-type: none"> • Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75 • Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80 • <i>Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> • <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i> • <i>Configuring Gigabit Ethernet Interfaces for EX Series Switches with ELS support</i> • <i>Configuring Gigabit Ethernet Interfaces (J-Web Procedure)</i> • Junos OS Ethernet Interfaces Configuration Guide |

lacp (802.3ad)

| | |
|-------------------------------------|---|
| Syntax | <pre>lacp { force-up; (primary backup); port-priority; }</pre> |
| Hierarchy Level (EX Series) | <p>[edit interfaces <i>interface-name</i> ether-options 802.3ad]</p> <p>[edit interfaces aeX aggregated-ether-options]</p> <p>[edit chassis aggregated-devices ethernet]</p> |
| Hierarchy Level (QFX Series) | <p>[edit interfaces <i>interface-name</i> ether-options 802.3ad]</p> |
| Release Information | <p>Statement introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Support for LACP link protection introduced in Junos OS Release 11.4 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>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.</p> <p>Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.</p> |
| | <p> NOTE: The port-priority statement is not supported on QFabric systems.</p> |
| | <p> NOTE: The force-up statement is not supported on QFX10002 switches.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> 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)*
- *Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80*
- *Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85*
- *Junos OS Ethernet Interfaces Configuration Guide*
- *Configuring Link Aggregation on page 81*
- *Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80*
- *Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75*

lacp (Aggregated Ethernet)

List of Syntax [Syntax \(NFX Series\) on page 248](#)
[Syntax \(EX Series\) on page 248](#)

Syntax (NFX Series)

```
lacp (active | passive) {
  admin-key key;
  fast-failover;
  link-protection {
    disable;
    (revertive | non-revertive);
  }
  periodic interval
  system-ID mac-address;
  system-priority priority;
  force-up;
}
```

Syntax (EX Series)

```
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 (EX Series)

[edit interfaces aeX aggregated-ether-options]
 [edit logical-systems *logical-system-name* interfaces aeX aggregated-ether-options]

Hierarchy Level (NFX Series)

[edit [interfaces interface-name](#) aggregated-ether-options]

Release Information

Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

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

For EX Series, when you configure the **accept-data** statement at the **[edit interfaces aeX 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 aeX 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.



NOTE: The `force-up` statement is not supported on QFX10002 switches.

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 Link Aggregation on page 81](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 80](#)
- [Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75](#)
- *Configuring LACP for Aggregated Ethernet Interfaces*

link-protection

| | |
|---------------------------------|--|
| Syntax | <pre>link-protection { disable; (revertive non-revertive); }</pre> |
| Hierarchy Level | <pre>[edit interfaces aex aggregated-ether-options] [edit interfaces aex aggregated-ether-options lACP]</pre> |
| Release Information | <p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 15.1F4 for PTX Series routers.</p> <p>Support for disable, revertive, and non-revertive statements added in Junos OS Release 9.3.</p> |
| Description | <p>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 gige-ether-options 802.3ad aex] hierarchy level or the [edit interfaces fe-fpc/pic/port fastether-options 802.3ad aex] hierarchy level.</p> <p>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.</p> <p>For Junos OS link protection, specify link-protection at the following hierarchy levels:</p> <ul style="list-style-type: none"> • [edit interfaces ge-fpc/pic/port ether-options 802.3ad aex] • [edit interfaces xe-fpc/pic/port ether-options 802.3ad aex] hierarchy level or at the [edit interfaces xe-fpc/pic/port ether-options 802.3ad aex] hierarchy level. <p>To disable link protection, use the delete interface ae aggregate-ether-options link-protection statement at the [edit interfaces aex aggregated-ether-options] hierarchy level or the [edit interfaces aex aggregated-ether-options lACP] hierarchy level.</p> |
| Options | The remaining statements are explained separately. See CLI Explorer . |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Configuring Aggregated Ethernet Link Protection</i> • <i>Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85</i> |

periodic

| | |
|-------------------------------------|--|
| List of Syntax | Syntax (EX Series) on page 252 Syntax (QFX Series) on page 252 |
| Syntax (EX Series) | <code>periodic <i>interval</i>;</code> |
| Syntax (QFX Series) | <code>periodic (fast slow);</code> |
| Hierarchy Level (EX Series) | [edit interfaces aex aggregated-ether-options lACP], [edit interfaces interface-range <i>name</i> aggregated-ether-options lACP] |
| Hierarchy Level (QFX Series) | [edit interfaces aex 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 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. 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 | <p><i>interval</i>—Interval for periodic transmission of LACP packets.</p> <ul style="list-style-type: none"> fast—Transmit packets every second. slow—Transmit packets every 30 seconds. <p>Default: fast</p> |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> Configuring LACP for Aggregated Ethernet Interfaces Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80 Example: Configuring Aggregated Ethernet High-Speed Uplinks Between an EX4200 Virtual Chassis Access Switch and an EX4200 Virtual Chassis Distribution Switch Configuring Aggregated Ethernet LACP (CLI Procedure) on page 80 Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75 Junos OS Network Interfaces Library for Routing Devices |

CHAPTER 11

Redundant Trunk Groups Configuration Statements

- [group \(Redundant Trunk Groups\) on page 254](#)
- [interface \(Redundant Trunk Groups\) on page 255](#)
- [preempt-cutover-timer on page 256](#)
- [redundant-trunk-group on page 257](#)

group (Redundant Trunk Groups)

| | |
|---------------------------------|--|
| Syntax | <pre>group <i>name</i> { interface <i>interface-name</i> <primary>; interface <i>interface-name</i>; preempt-cutover-timer <i>seconds</i>; }</pre> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with ELS: <pre>[edit switch-options redundant-trunk-group]</pre> For platforms without ELS: <pre>[edit ethernet-switching-options redundant-trunk-group]</pre> |
| Release Information | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Using the Enhanced Layer 2 Software CLI</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| Description | Create a redundant trunk group. |
| Options | <p>name—The name of the redundant trunk group.</p> <ul style="list-style-type: none"> For platforms with ELS: <p>The group name must be a string “rtgn” where n is a number from 0 through 15, such as “rtg2” or “rtg10”.</p> For platforms without ELS: <p>The group name must start with a letter and can consist of letters, numbers, dashes, and underscores.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches</i> <i>Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support</i> <i>Understanding Redundant Trunk Links (Legacy RTG Configuration)</i> |

interface (Redundant Trunk Groups)

| | |
|---------------------------------|---|
| Syntax | <pre>interface <i>interface-name</i> <primary>; interface <i>interface-name</i>;</pre> |
| Hierarchy Level | <p>For platforms with ELS:</p> <pre>[edit switch-options redundant-trunk-group <i>group name</i>]</pre> <p>For platforms without ELS:</p> <pre>[edit ethernet-switching-options redundant-trunk-group <i>group name</i>]</pre> |
| Release Information | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Using the Enhanced Layer 2 Software CLI</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| Description | <p>Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over as the primary link without waiting for normal STP convergence.</p> |
| Options | <p>interface <i>interface-name</i>—A logical interface or an aggregated interface containing multiple ports.</p> <p>primary—(Optional) Specify one of the interfaces in the redundant group as the primary link. The interface without this option is the secondary link in the redundant group. If a link is not specified as primary, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are ge-0/1/0 and ge-0/1/1, the software assigns ge-0/1/1 as the active link.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system—control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches</i> • <i>Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support</i> • <i>Understanding Redundant Trunk Links (Legacy RTG Configuration)</i> |

preempt-cutover-timer

| | |
|---------------------------------|--|
| Syntax | <code>preempt-cutover-timer seconds;</code> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with ELS: <pre>[edit switch-options redundant-trunk-group group name] [edit interfaces name aggregated-ether-options lacp link-protection rtg-config] [edit interfaces name aggregated-ether-options link-protection rtg-config]</pre> For platforms without ELS: <pre>[edit ethernet-switching-options redundant-trunk-group group name]</pre> |
| Release Information | <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Using the Enhanced Layer 2 Software CLI</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| Description | Change the length of time that a re-enabled primary link waits to take over from an active secondary link in a redundant trunk group (RTG). |
| Default | If you do not change the time with the preempt-cutover-timer statement, a re-enabled primary link takes over from the active secondary link after 1 second. |
| Options | <p>seconds—Number of seconds that the primary link waits to take over from the active secondary link.</p> <p>Range: 1 through 600 seconds</p> |
| Required Privilege Level | <p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches</i> <i>Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support</i> <i>Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection</i> |

redundant-trunk-group

| | |
|---------------------------------|---|
| Syntax | <pre> redundant-trunk-group { group name { interface interface-name <primary>; interface interface-name; preempt-cutover-timer seconds; } } </pre> |
| Hierarchy Level | <ul style="list-style-type: none"> For platforms with ELS: <div>[edit switch-options]</div> For platforms without ELS: <div>[edit ethernet-switching-options]</div> |
| Release Information | <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Using the Enhanced Layer 2 Software CLI</i> for information about ELS.)</p> <p>Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> |
| Description | <p>Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal spanning-tree protocol convergence.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p> |
| Required Privilege Level | <p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> <i>Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches</i> <i>Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support</i> <i>Understanding Redundant Trunk Links (Legacy RTG Configuration)</i> |

CHAPTER 12

Ethernet OAM Link Fault Management Operational Command

- `show oam ethernet link-fault-management`

show oam ethernet link-fault-management

Syntax `show oam ethernet link-fault-management
<brief | detail>
<interface-name>`

Release Information Command introduced in Junos OS Release 9.4 for EX Series switches.

Description Displays Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.

Options **brief | detail**—(Optional) Display the specified level of output.

interface-name —(Optional) Display link fault management information for the specified Ethernet interface only.

Required Privilege Level view

Related Documentation

- *Example: Configuring Ethernet OAM Link Fault Management*
- [Configuring Ethernet OAM Link Fault Management on page 4](#)

List of Sample Output [show oam ethernet link-fault-management brief on page 264](#)
[show oam ethernet link-fault-management detail on page 264](#)

Output Fields [Table 25 on page 260](#) lists the output fields for the **show oam ethernet link-fault-management** command. Output fields are listed in the approximate order in which they appear.

Table 25: show oam ethernet link-fault-management Output Fields

| Field Name | Field Description | Level of Output |
|------------------------|---|-----------------|
| Status | Indicates the status of the established link. <ul style="list-style-type: none"> • Fail—A link fault condition exists. • Running—A link fault condition does not exist. | All levels |
| Discovery state | State of the discovery mechanism: <ul style="list-style-type: none"> • Passive Wait • Send Any • Send Local Remote • Send Local Remote Ok | All levels |
| Peer address | Address of the OAM peer. | All levels |

Table 25: show oam ethernet link-fault-management Output Fields (continued)

| Field Name | Field Description | Level of Output |
|----------------------------------|--|-----------------|
| Flags | Information about the interface. <ul style="list-style-type: none"> • Remote-Stable—Indicates remote OAM client acknowledgment of, and satisfaction with local OAM state information. False indicates that remote DTE has either not seen or is unsatisfied with local state information. True indicates that remote DTE has seen and is satisfied with local state information. • Local-Stable—Indicates local OAM client acknowledgment of, and satisfaction with remote OAM state information. False indicates that local DTE either has not seen or is unsatisfied with remote state information. True indicates that local DTE has seen and is satisfied with remote state information. • Remote-State-Valid—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. False indicates that OAM client has not seen remote state information. True indicates that the OAM client has seen remote state information. | All levels |
| Remote loopback status | Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session). | All levels |
| Remote entity information | Remote entity information. <ul style="list-style-type: none"> • Remote MUX action—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • Remote parser action—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • Discovery mode—Indicates whether discovery mode is active or inactive. • Unidirectional mode—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes. • Remote loopback mode—Indicates whether remote loopback is supported or not supported. • Link events—Indicates whether interpreting link events is supported or not supported on the remote peer. • Variable requests—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer. | All levels |
| OAM Receive Statistics | | |
| Information | The number of information PDUs received. | detail |
| Event | The number of loopback control PDUs received. | detail |
| Variable request | The number of variable request PDUs received. | detail |
| Variable response | The number of variable response PDUs received. | detail |
| Loopback control | The number of loopback control PDUs received. | detail |

Table 25: show oam ethernet link-fault-management Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--|--|-----------------|
| Organization specific | The number of vendor organization specific PDUs received. | detail |
| OAM Transmit Statistics | | |
| Information | The number of information PDUs transmitted. | detail |
| Event | The number of event notification PDUs transmitted. | detail |
| Variable request | The number of variable request PDUs transmitted. | detail |
| Variable response | The number of variable response PDUs transmitted. | detail |
| Loopback control | The number of loopback control PDUs transmitted. | detail |
| Organization specific | The number of vendor organization specific PDUs transmitted. | detail |
| OAM Received Symbol Error Event information | | |
| Events | The number of symbol error event TLVs that have been received after the OAM sublayer was reset. | detail |
| Window | The symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer. | detail |
| Threshold | The number of errored symbols in the period required for the event to be generated. | detail |
| Errors in period | The number of symbol errors in the period reported in the received event PDU. | detail |
| Total errors | The number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors. | detail |
| OAM Received Frame Error Event Information | | |
| Events | The number of errored frame event TLVs that have been received after the OAM sublayer was reset. | detail |
| Window | The duration of the window in terms of the number of 100 ms period intervals. | detail |
| Threshold | The number of detected errored frames required for the event to be generated. | detail |
| Errors in period | The number of detected errored frames in the period. | detail |

Table 25: show oam ethernet link-fault-management Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--|---|-----------------|
| Total errors | The number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer. | detail |
| OAM Received Frame Period Error Event Information | | |
| Events | The number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset. | detail |
| Window | The duration of the frame seconds window. | detail |
| Threshold | The number of frame seconds errors in the period. | detail |
| Errors in period | The number of frame seconds errors in the period. | detail |
| Total errors | The number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset. | detail |
| OAM Transmitted Symbol Error Event Information | | |
| Events | The number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset. | detail |
| Window | The symbol error event window in the transmitted PDU. | detail |
| Threshold | The number of errored symbols in the period required for the event to be generated. | detail |
| Errors in period | The number of symbol errors in the period reported in the transmitted event PDU. | detail |
| Total errors | The number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset. | detail |
| OAM Transmitted Frame Error Event Information | | |
| Events | The number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset. | detail |
| Window | The duration of the window in terms of the number of 100 ms period intervals. | detail |
| Threshold | The number of detected errored frames required for the event to be generated. | detail |
| Errors in period | The number of detected errored frames in the period. | detail |
| Total errors | The number of errored frames that have been detected after the OAM sublayer was reset. | detail |

Sample Output

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
```

```
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:72:2c:83
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
  Remote MUX action: discarding, Remote parser action: loopback
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
```

```
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:0a:07:14
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
  Information: 186365, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM transmit statistics:
  Information: 186347, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM received symbol error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame period error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

CHAPTER 13

Interfaces Operational Commands

- `monitor interface`
- `show interfaces diagnostics optics`
- `show interfaces ge`
- `show interfaces irb`
- `show interfaces xe`

monitor interface

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

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

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



NOTE: On Junos OS Evolved, you can use the `monitor interface` command over SSH sessions, but console and Telnet sessions are not supported.



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 26 on page 267](#). The keys are not case-sensitive.

Table 26: 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. |
| 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 27 on page 267](#). The keys are not case-sensitive.

Table 27: 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 269](#)
[monitor interface \(OTN Interface\) on page 270](#)
[monitor interface \(MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface\) on page 271](#)
[monitor interface \(MX480 Router with MPC5E and 100-Gigabit Ethernet Interface\) on page 272](#)

[monitor interface \(MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 273](#)
[monitor interface \(MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface\) on page 274](#)
[monitor interface \(MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 274](#)
[monitor interface \(Logical\) on page 275](#)
[monitor interface \(QFX3500 Switch\) on page 275](#)
[monitor interface traffic on page 276](#)
[monitor interface traffic \(QFX3500 Switch\) on page 276](#)
[monitor interface traffic detail \(QFX3500 Switch\) on page 277](#)

Output Fields [Table 28 on page 268](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 28: monitor interface Output Fields

| Field Name | Field Description | Level of Output |
|-------------------------|--|-----------------|
| routerl | 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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 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 28: monitor interface Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------|--|-----------------|
| Remote Statistics | <p>(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.</p> <ul style="list-style-type: none"> • 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 | <p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> • 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 <i>interface-name</i>] hierarchy level. | detail |

Sample Output

monitor interface (Physical)

```

user@host> monitor interface so-0/0/0

router1                               Seconds: 19                               Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: 0C48
Traffic statistics:
Input packets:                        6045 (0 pps)                               [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:
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:
  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:
    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
    BIP                             0                  [0]
    BBE                             0                  [0]
    ES                              24                  [0]
    SES                             0                  [0]
    UAS                             1255               [0]
OTN ODU - Counters

```

```

BIP                                0
BBE                                0                [0]
ES                                 24                [0]
SES                                0                [0]
UAS                                1256             [0]
OTN ODU - Received Overhead       0                [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:
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: 1
Output errors: 0
Output drops: 0
Aged packets: 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:
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: 3
Output errors: 0
Output drops: 0
Aged packets: 0
OTN Link 0
OTN Alarms:
OTN Defects:
OTN OC - Seconds
LOS 0
LOF 0

```

```

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                     0
OTN ODU - Counters
  BIP                     0
  BBE                     0
  ES                      0
  SES                     0
  UAS                     0
OTN ODU - Received Overhead
  APSPCC 0-3:            00 00 00 00

```

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

monitor interface (Logical)

```

user@host> monitor interface so-1/0/0.0

host name          Seconds: 16          Time: 15:33:39
                                   Delay: 0/0/1

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

```

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

monitor interface (QFX3500 Switch)

```

user@switch> monitor interface ge-0/0/0

Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:
  Input bytes:      0 (0 bps)
  Output bytes:     0 (0 bps)
  Input packets:    0 (0 pps)
  Output packets:   0 (0 pps)
Error statistics:
  Input errors:     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 | |

show interfaces diagnostics optics

| | |
|---------------------------------|---|
| Syntax | <code>show interfaces diagnostics optics <i>interface-name</i></code> |
| Release Information | <p>Command introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP, SFP+, XFP, QSFP+, or CFP) installed in EX Series or QFX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p> |
| Options | <i>interface-name</i> —Name of the interface associated with the port in which the transceiver is installed: <i>ge-fpc/pic/port</i> , <i>xe-fpc/pic/port</i> , or <i>et-fpc/pic/port</i> . |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • <i>Monitoring Interface Status and Traffic</i> • Monitoring Interface Status and Traffic on page 65 • <i>Installing a Transceiver</i> • <i>Installing a Transceiver in a QFX Series Device</i> • <i>Removing a Transceiver</i> • <i>Removing a Transceiver from a QFX Series Device</i> • Junos OS Ethernet Interfaces Configuration Guide |
| List of Sample Output | <p>show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 286</p> <p>show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 287</p> <p>show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 288</p> <p>show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver) on page 289</p> <p>show interfaces diagnostics optics et-4/1/0 (CFP Transceiver) on page 290</p> |
| Output Fields | Table 29 on page 280 lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the approximate order in which they appear. |

Table 29: show interfaces diagnostics optics Output Fields

| Field Name | Field Description |
|--|--|
| Physical interface | Displays the name of the physical interface. |
| Lane 'x' | Displays an individual Tx/Rx data transmission channel "x" associated with the given physical interface device (or IFD). |
| 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 . |

Table 29: show interfaces diagnostics optics Output Fields (continued)

| Field Name | Field Description |
|---|---|
| Laser output power high warning (Not available for QSFP+ transceivers) | Displays whether the laser output power high warning is On or Off . |
| 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 . |

Table 29: show interfaces diagnostics optics Output Fields (continued)

| Field Name | Field Description |
|---|--|
| Module voltage low warning (Not available for XFP and QSFP+ transceivers) | Displays whether the module voltage low warning is On or Off . |
| 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 . |

Table 29: show interfaces diagnostics optics Output Fields (continued)

| Field Name | Field Description |
|--|--|
| Tx laser disabled alarm (Not available for SFP, SFP+, XFP, and QSFP+ transceivers) | Displays whether the Tx laser disabled alarm is On or Off . |
| 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. |

Table 29: show interfaces diagnostics optics Output Fields (continued)

| Field Name | Field Description |
|--|--|
| Laser output power low alarm threshold (Not available for QSFP+ transceivers) | Displays the vendor-specified threshold for the laser output power low alarm. |
| 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. |

Table 29: show interfaces diagnostics optics Output Fields (continued)

| Field Name | Field Description |
|---|--|
| Laser rx power high warning threshold (Not available for QSFP+ transceivers) | Displays the vendor-specified threshold for the laser rx power high warning. |
| 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 . |

Table 29: show interfaces diagnostics optics Output Fields (continued)

| Field Name | Field Description |
|---|--|
| 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 . |
| 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

List of Syntax [Syntax \(EX Series\) on page 294](#)
[Syntax \(QFX Series\) on page 294](#)

Syntax (EX Series) `show interfaces ge-fpc/pic/port`
`<brief | detail | extensive | terse>`
`<media>`
`<statistics>`

Syntax (QFX Series) `show interfaces device-name:type-fpc/pic/port`
`<brief | detail | extensive | terse>`
`<descriptions>`
`<media>`
`<routing-instance (all | instance-name)>`
`<snmp-index snmp-index>`
`<statistics>`

Release Information Command introduced in Junos OS Release 9.0 for EX Series switches.
 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 status information about the specified Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.



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***—(EX Series) Display standard information about the specified Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) (EX Series) Display the specified level of output.

media—(Optional) (EX Series) Display media-specific information about network interfaces.

statistics—(Optional) (EX Series) Display static interface statistics.

brief | detail | extensive | terse—(Optional) (QFX Series) Display the specified level of output.

device-name:type-fpc/pic/port— (QFX Series) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.

descriptions—(Optional) (QFX Series) Display interface description strings.

media—(Optional) (QFX Series) Display media-specific information about network interfaces.

routing instance (all | *instance-name*)—(Optional) (QFX Series) Display the name of an individual routing-instance or display all routing-instances.

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

statistics—(Optional) (QFX Series) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- *Monitoring Interface Status and Traffic*
- *Troubleshooting Network Interfaces on EX3200 Switches*
- *Troubleshooting Network Interfaces on EX4200 Switches*
- *Troubleshooting an Aggregated Ethernet Interface*
- [Junos OS Ethernet Interfaces Configuration Guide](#)
- [Monitoring Interface Status and Traffic on page 65](#)
- [Troubleshooting Network Interfaces on page 65](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 92](#)
- *Junos OS Network Interfaces Library for Routing Devices*

List of Sample Output

[show interfaces ge-0/0/0 on page 302](#)
[show interfaces ge-0/0/0 brief on page 302](#)
[show interfaces ge-0/0/0 brief \(with EEE Enabled on the EEE-capable Base-T copper Ethernet interfaces\) on page 303](#)
[show interfaces ge-0/0/0 detail on page 303](#)
[show interfaces ge-0/0/4 extensive on page 304](#)

Output Fields

[Table 30 on page 295](#) lists the output fields for the **show interfaces ge-** command. Output fields are listed in the approximate order in which they appear.

Table 30: 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 |

Table 30: show interfaces ge- 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 |
| 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: <ul style="list-style-type: none"> • 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 |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Current address | Configured MAC address. | detail extensive none |

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

| Field Name | Field Description | Level of Output |
|-------------------------|---|-----------------------|
| 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 | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • 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. <p>NOTE: The bandwidth bps counter is not enabled on the switch.</p> | detail extensive |
| Input errors | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 30: show interfaces ge- Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---|--|------------------------------|
| Output errors | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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) | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • 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 | <p>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.</p> <ul style="list-style-type: none"> • 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 30: *show interfaces ge-* Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------------|---|------------------|
| MAC statistics | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • 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 30: *show interfaces ge- Output Fields (continued)*

| Field Name | Field Description | Level of Output |
|--|---|-----------------|
| Autonegotiation information | <p>Information about link autonegotiation:</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 | <p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number: <ul style="list-style-type: none"> • 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 30: *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. <i>NOTE:</i> 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 |
| <i>protocol-family</i> | 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 30: *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 interface. | 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 IEEE Enabled on the IEEE-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 | <pre>show interfaces irb <brief detail extensive terse> <descriptions> <media> <routing-instance <i>instance-name</i>> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | <p>Command introduced in Junos OS Release 12.3R2.</p> <p>Command introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series</p> |
| Description | Display integrated routing and bridging interfaces information. |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>routing-instance <i>instance-name</i>—(Optional) Display information for the interface with the specified SNMP index.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the interface with the specified SNMP index.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| 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 | <p>show interfaces irb extensive on page 311</p> <p>show interfaces irb snmp-index on page 313</p> |
| Output Fields | Table 31 on page 307 lists the output fields for the show interfaces irb command. Output fields are listed in the approximate order in which they appear. |

Table 31: show interfaces irb Output Fields

| Field Name | Field Description | Level of Output |
|--------------------|-------------------|-----------------|
| Physical Interface | | |

Table 31: show interfaces irb Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------------|--|------------------------------------|
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the physical interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> . | 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 <i>Common Output Fields Description</i> . | detail extensive brief none |
| Interface flags | Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> . | 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 <i>Common Output Fields Description</i> . | 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 |

Table 31: show interfaces irb Output Fields (continued)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| 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 |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • 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 | <p>Number of IPv6 transit bytes and packets received and transmitted on the physical interface if IPv6 statistics tracking is enabled.</p> <ul style="list-style-type: none"> • 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 | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 |

Table 31: show interfaces irb Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------------------|---|---------------------------------|
| Output errors | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 |
| 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 <i>Common Output Fields Description</i> . | detail extensive |
| Encapsulation | Encapsulation on the logical interface. | detail extensive |
| Bandwidth | Dummy value that is ignored by an IRB interface. IRB interfaces are pseudo interfaces and do not have physical bandwidth associated with them. | detail extensive |
| Routing Instance | Routing instance IRB is configured under. | detail extensive |
| Bridging Domain | Bridging domain IRB is participating in. | detail extensive |
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the logical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface • Output packets—Number of packets transmitted on the interface. | detail extensive |

Table 31: show interfaces irb Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------|---|--------------------------|
| IPv6 transit statistics | Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled. <ul style="list-style-type: none"> 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 <i>Common Output Fields Description</i> . | detail extensive |
| MTU | Maximum transmission unit size on the logical interface. | detail extensive |
| 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 <i>Common Output Fields Description</i> . | 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 <i>Common Output Fields Description</i> . | 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, Runt: 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 xe

List of Syntax [Syntax \(QFX Series\) on page 314](#)
[Syntax \(EX Series\) on page 314](#)

Syntax (QFX Series) `show interfaces device-name:type-fpc/pic/port`
`<brief | detail | extensive | terse>`
`<descriptions>`
`<media>`
`<routing-instance (all | instance-name)>`
`<snmp-index snmp-index>`
`<statistics>`

Syntax (EX Series) `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.
 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 status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces.



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 ***device-name:type-fpc/pic/port***—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.

brief | detail | extensive | terse—(Optional) (QFX Series) Display the specified level of output.

descriptions—(Optional) (QFX Series) Display interface description strings.

media—(Optional) (QFX Series) Display media-specific information about network interfaces.

routing-instance (all | *instance-name*)—(Optional) (QFX Series) Display the name of an individual routing instance or display all routing instances.

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

statistics—(Optional) (QFX Series) Display static interface statistics.

xe-fpc/pic/port —(EX Series) Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) (EX Series) Display the specified level of output.

media—(Optional) (EX Series) 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) (EX Series) 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 65](#)
- [Monitoring Interface Status and Traffic](#)
- [Troubleshooting Network Interfaces on page 65](#)
- [Troubleshooting an Aggregated Ethernet Interface on page 92](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)
- [Troubleshooting Network Interfaces on EX3200 Switches](#)
- [Troubleshooting Network Interfaces on EX4200 Switches](#)
- [Troubleshooting an Aggregated Ethernet Interface](#)
- [Junos OS Ethernet Interfaces Configuration Guide](#)

List of Sample Output

[show interfaces on page 323](#)
[show interfaces \(Asymmetric Flow Control\) on page 324](#)
[show interfaces brief on page 324](#)
[show interfaces detail on page 324](#)

[show interfaces detail \(Asymmetric Flow Control\) on page 326](#)
[show interfaces extensive on page 327](#)
[show interfaces extensive \(Asymmetric Flow Control\) on page 329](#)
[show interfaces terse on page 332](#)
[show interfaces \(QFabric System\) on page 332](#)

Output Fields Table 32 on page 316 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

Table 32: show interfaces xe 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. | 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 |
| Link-level type | Encapsulation being used on the physical interface. | All levels |
| MTU | Maximum transmission unit size on the physical interface. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Duplex | Duplex mode of the interface, either Full-Duplex or Half-Duplex . | 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 |
| LAN-PHY mode | 10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications. | All levels |
| Unidirectional | Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces. | All levels |
| Flow control | Flow control status: Enabled or Disabled . | All levels |
| NOTE: This field is only displayed if asymmetric flow control is not configured. | | |

Table 32: *show interfaces xe Output Fields (continued)*

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| Configured-flow-control | Configured flow control for the interface transmit buffers (tx-buffers) and receive buffers (rx-buffers): <ul style="list-style-type: none"> tx-buffers—On if the interface is configured to respond to Ethernet PAUSE messages received from the connected peer. Off if the interface is not configured to respond to received PAUSE messages. rx-buffers—On if the interface is configured to generate and send Ethernet PAUSE messages to the connected peer. Off if the interface is not configured to generate and send PAUSE messages. <p>NOTE: This field is only displayed if asymmetric flow control is configured.</p> | All levels |
| Auto-negotiation | Autonegotiation status: Enabled or Disabled . | All levels |
| Remote-fault | Remote fault status: <ul style="list-style-type: none"> Online—Autonegotiation is manually configured as online. Offline—Autonegotiation is manually configured as offline. | 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 |
| Wavelength | Configured wavelength, in nanometers (nm). | All levels |
| Frequency | Frequency associated with the configured wavelength, in terahertz (THz). | All levels |
| CoS queues | Number of CoS queues configured. | detail extensive none |
| Schedulers | Number of CoS schedulers configured. | 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 | Hardware MAC address. | 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 |
| Input Rate | Input rate in bits per second (bps) and packets per second (pps). | None specified |
| Output Rate | Output rate in bps and pps. | None specified |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |

Table 32: *show interfaces xe Output Fields (continued)*

| Field Name | Field Description | Level of Output |
|---------------------------|---|-------------------------|
| Traffic statistics | <p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • 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. <p>NOTE: The bandwidth bps counter is not enabled.</p> | detail extensive |
| Input errors | <p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 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 32: *show interfaces xe Output Fields (continued)*

| Field Name | Field Description | Level of Output |
|---------------------------------|--|-------------------------|
| Output errors | <p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • 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 router 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) | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • 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 |
| Queue Number | The CoS queue number and the forwarding classes mapped to the queue number. The Mapped forwarding class column lists the forwarding classes mapped to each CoS queue. | detail extensive |
| Ingress queues | Total number of ingress queues supported on the specified interface. | extensive |
| Queue counters (Ingress) | <p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. | extensive |

Table 32: *show interfaces xe Output Fields (continued)*

| Field Name | Field Description | Level of Output |
|---|--|------------------------------|
| Active alarms and Active defects | <p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • 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 |
| PCS statistics | Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device. | detail extensive |
| MAC statistics | <p>Receive and Transmit statistics reported by the PIC's MAC subsystem.</p> <ul style="list-style-type: none"> • 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 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 packets that exceeds the configured MTU. • 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 runs (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. This counter is not supported on EX Series switches and is always displayed as 0. • 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 32: *show interfaces xe Output Fields (continued)*

| Field Name | Field Description | Level of Output |
|-----------------------------|--|-----------------|
| Autonegotiation information | <p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when the Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE 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. • Local resolution: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are enable or disable. • 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 |

Table 32: *show interfaces xe Output Fields (continued)*

| Field Name | Field Description | Level of Output |
|---|--|------------------------------|
| Packet Forwarding Engine configuration | Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> Destination slot—FPC slot number. CoS transmit queue—Queue number and its associated user-configured forwarding class name. Bandwidth %—Percentage of bandwidth allocated to the queue. Bandwidth bps—Bandwidth allocated to the queue (in bps). Buffer %—Percentage of buffer space allocated to the queue. Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. 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 |
| 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 |
| Protocol | Protocol family. | detail extensive none |
| 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 statics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface. | 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 |

Table 32: *show interfaces xe* Output Fields (continued)

| Field Name | Field Description | Level of Output |
|-------------------------|--|------------------------------|
| 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 (uRPF) is explicitly configured on the specified interface, the uRPF flag appears. If uRPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag does not appear even though uRPF is enabled. | detail extensive |
| Addresses, Flags | Information about the address flags. | detail extensive none |
| <i>protocol-family</i> | 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 flag. | 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

```
user@switch> show interfaces xe-0/0/1
```

```
Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
  Flow control: Disabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped  : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects: None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
```

```

Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Input packets : 0
Output packets: 0
Protocol eth-switch, MTU: 0
Flags: Trunk-Mode

```

show interfaces (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1
```

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Configured-flow-control tx-buffers: off rx-buffers: on
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 12 supported, 12 maximum usable queues
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped   : 2011-06-01 00:42:03 PDT (00:02:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523)
  Flags: SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol eth-switch, MTU: 0
  Flags: Trunk-Mode

```

show interfaces brief

```
user@switch> show interfaces xe-0/0/1 brief
```

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

Logical interface xe-0/0/1.0
  Flags: SNMP-Traps Encapsulation: ENET2
  eth-switch

```

show interfaces detail

```
user@switch> show interfaces xe-0/0/1 detail
```

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:

```

```

Disabled,
Flow control: Disabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags     : None
CoS queues     : 12 supported, 12 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped   : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps
Input packets : 0          0 pps
Output packets: 0          0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort          0              0              0
  1 fc7                  0              0              0
  2 no-loss              0              0              0
  3 fcoe                 0              0              0
  4 fc4                  0              0              0
  5 fc5                  0              0              0
  6 fc6                  0              0              0
  7 network-cont         0              0              0
  8 mcast                0              0              0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast
Active alarms   : None
Active defects  : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0

```

```

Output packets:                0
Local statistics:
Input bytes :                  0
Output bytes :                  0
Input packets:                 0
Output packets:                 0
Transit statistics:
Input bytes :                  0          0 bps
Output bytes :                  0          0 bps
Input packets:                 0          0 pps
Output packets:                 0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces detail (Asymmetric Flow Control)

user@switch> show interfaces xe-0/0/1 detail

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:02:50 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:14 ago)
Traffic statistics:
Input bytes :                  0          0 bps
Output bytes :                  0          0 bps
Input packets:                 0          0 pps
Output packets:                 0          0 pps
IPv6 transit statistics:
Input bytes :                  0
Output bytes :                  0
Input packets:                 0
Output packets:                 0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets

0 best-effort                0                0                0
1 fc7                        0                0                0
2 no-loss                    0                0                0
3 fcoe                        0                0                0
4 fc4                         0                0                0
5 fc5                         0                0                0
6 fc6                         0                0                0
7 network-cont                0                0                0

```



```

      8 mcast                                0                                0                                0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast
Active alarms : None
Active defects : None

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces extensive

```
user@switch> show interfaces xe-0/0/1 extensive
```

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Disabled
  Device flags : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags : None
  CoS queues : 12 supported, 12 maximum usable queues
  Hold-times : Up 0 ms, Down 0 ms
  Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
  Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
  Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:

```

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort        0                0                0
1 fc7                0                0                0
2 no-loss            0                0                0
3 fcoe              0                0                0
4 fc4               0                0                0
5 fc5               0                0                0
6 fc6               0                0                0
7 network-cont      0                0                0
8 mcast             0                0                0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets    0            0
Unicast packets  0            0
Broadcast packets 0            0
Multicast packets 0            0
CRC/Align errors 0            0
FIFO errors       0            0
MAC control frames 0            0
MAC pause frames  0            0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
MAC Priority Flow Control Statistics:

```

```

Priority : 0          0          0
Priority : 1          0          0
Priority : 2          0          0
Priority : 3          0          0
Priority : 4          0          0
Priority : 5          0          0
Priority : 6          0          0
Priority : 7          0          0
Filter statistics:
  Input packet count      0
  Input packet rejects    0
  Input DA rejects        0
  Input SA rejects        0
  Output packet count     0
  Output packet pad count 0
  Output packet error count 0
  CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
    0 best-effort         75      7500000000    75          0      low
none
    7 network-control     5       500000000    5           0      low
none
    8 mcast                20     2000000000    20          0      low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes :          0
  Output bytes :          0
  Input packets:          0
  Output packets:          0
Local statistics:
  Input bytes :          0
  Output bytes :          0
  Input packets:          0
  Output packets:          0
Transit statistics:
  Input bytes :          0          0 bps
  Output bytes :          0          0 bps
  Input packets:          0          0 pps
  Output packets:          0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces extensive (Asymmetric Flow Control)

```
user@switch> show interfaces xe-0/0/1 extensive
```

```

Physical interface: xe-0/0/1, Enabled, Physical link is Up
  Interface index: 49195, SNMP ifIndex: 591, Generation: 169
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
  Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,

```

```

Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters: Queued packets Transmitted packets Dropped packets

0 best-effort 0 0 0
1 fc7 0 0 0
2 no-loss 0 0 0
3 fcoe 0 0 0
4 fc4 0 0 0
5 fc5 0 0 0
6 fc6 0 0 0
7 network-cont 0 0 0
8 mcast 0 0 0

Queue number: Mapped forwarding classes
0 best-effort
1 fc7
2 no-loss
3 fcoe
4 fc4
5 fc5
6 fc6
7 network-control
8 mcast
Active alarms : None
Active defects : None
MAC statistics: Receive Transmit

```

```

Total octets                0                0
Total packets               0                0
Unicast packets             0                0
Broadcast packets          0                0
Multicast packets          0                0
CRC/Align errors           0                0
FIFO errors                 0                0
MAC control frames         0                0
MAC pause frames           0                0
Oversized frames           0
Jabber frames              0
Fragment frames            0
VLAN tagged frames         0
Code violations             0
MAC Priority Flow Control Statistics:
  Priority : 0              0                0
  Priority : 1              0                0
  Priority : 2              0                0
  Priority : 3              0                0
  Priority : 4              0                0
  Priority : 5              0                0
  Priority : 6              0                0
  Priority : 7              0                0
Filter statistics:
  Input packet count        0
  Input packet rejects      0
  Input DA rejects          0
  Input SA rejects          0
  Output packet count       0
  Output packet pad count   0
  Output packet error count 0
  CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue
    %      Bandwidth      %      Buffer Priority Limit
    %      bps            %      usec
0 best-effort    75    7500000000    75    0      low  none
7 network-control 5     500000000    5     0      low  none
8 mcast         20    2000000000    20    0      low  none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0                0 bps
  Output bytes : 0                0 bps
  Input packets: 0                0 pps
  Output packets: 0                0 pps

```

```
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode
```

show interfaces terse

```
user@switch> show interfaces xe-0/0/1 terse
```

| Interface | Admin | Link | Proto | Local | Remote |
|------------|-------|------|------------|-------|--------|
| xe-0/0/1 | up | up | | | |
| xe-0/0/1.0 | up | up | eth-switch | | |

show interfaces (QFabric System)

```
user@switch> show interfaces node1:xe-0/0/0
```

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

CHAPTER 14

LAGs and LACP Operational Commands

- `show lacp interfaces`
- `show lacp statistics interfaces (View)`

show lacp interfaces

Syntax `show lacp interfaces`
`<interface-name>`
`extensive`

Release Information Command introduced in Junos OS Release 7.6.
extensive statement introduced in Junos OS Release 16.1R1
 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, Fast 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—**aenumber**
- Fast Ethernet—**fe-fpc/pic/port**
- Gigabit Ethernet—**ge-fpc/pic/port**
- 10 Gigabit Ethernet—**xe-fpc/pic/port**

extensive—Display LACP information for the interface in detail.



NOTE: The `show lacp interfaces` command returns the following error message if your system is not configured in either active or passive LACP mode:

“Warning: lacp subsystem not running – not needed by configuration”

Required Privilege Level view

Related Documentation

- [Configuring Aggregated Ethernet Links \(CLI Procedure\)](#)
- [Configuring Link Aggregation on page 81](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 80](#)
- [Configuring Aggregated Ethernet LACP \(CLI Procedure\) on page 80](#)

- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces for Switches on page 85](#)
- [Understanding Aggregated Ethernet Interfaces and LACP for Switches on page 75](#)
- [Junos OS Interfaces Fundamentals Configuration Guide](#)

List of Sample Output [show lacp interfaces \(Aggregated Ethernet\) on page 337](#)
[show lacp interfaces \(Gigabit Ethernet\) on page 337](#)
[show lacp interfaces \(10 Gigabit Ethernet\) on page 338](#)

Output Fields [Table 33 on page 335](#) lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 33: show lacp interfaces Output Fields

| Field Name | Field Description | Level of Output |
|----------------------|--|-----------------|
| LACP State | For a child interface configured with the force-up statement, LACP state displays FUP along with the interface name. | All Levels |
| Aggregated interface | Aggregated interface value. | All Levels |
| LACP State | <p>LACP state information for each aggregated interface:</p> <ul style="list-style-type: none"> • Role—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> • Actor—Local device participating in LACP negotiation. • Partner—Remote device participating in LACP negotiation. • Exp—Expired state. Yes indicates the actor or partner is in an expired state. No indicates 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, administratively configured for the partner. No indicates the operational partner information in use has been received in an LACP PDU. • Dist—Distribution of outgoing frames. No indicates 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 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 synchronized. It 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. It is currently not in the right aggregation. • Aggr—Ability of 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 fast transmission rate, depending upon the expressed LACP timeout preference (Long Timeout or Short Timeout). • Activity—Actor 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. • Core isolation state down (CDN)— LACP interface state. Down indicates the LACP interface is down because all the eBGP sessions for Ethernet VPN (EVPN) are down. | All Levels |

Table 33: show lacp interfaces Output Fields (continued)

| Field Name | Field Description | Level of Output |
|---------------|--|-----------------|
| LACP Protocol | <p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> • Link state (active or standby) indicated in parentheses next to the interface when link protection is configured. • Receive State—One of the following values: <ul style="list-style-type: none"> • 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 state machine. One of the following values: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Attached—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 may be passed down from the aggregator's distribution function for transmission. No indicates the transmit function of this link is not enabled. • Waiting—Multiplexer state machine is in a holding process, awaiting an outcome. | All Levels |

Table 33: show lacp interfaces Output Fields (continued)

| Field Name | Field Description | Level of Output |
|------------|--|-----------------|
| LACP info | <ul style="list-style-type: none"> • Role can be one of the following: <ul style="list-style-type: none"> • Actor—Local device participating in LACP negotiation. • Partner—Remote device participating in LACP negotiation. • System priority—Priority assigned to the system (by management or administrative policy), encoded as an unsigned integer. • System identifier—Actor or partner system ID, encoded as a MAC address. • Port priority—Priority assigned to the port by the actor or partner (by management or administrative policy), encoded as an unsigned integer. • Port number—Port number assigned to the port by the actor or partner, encoded as an unsigned integer. • Port key—Operational key value assigned to the port by the actor or partner, encoded as an unsigned integer. | Extensive |

Sample Output

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0 extensive
```

| | | | | | | | | | |
|----------------|---------------|---------|----------|-------------------|------------|-----|-------------------------|---------|----------|
| LACP state: | Role | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
| ge-0/0/1 | Actor | No | No | Yes | Yes | Yes | Yes | Fast | Active |
| ge-0/0/1 | Partner | No | No | Yes | Yes | Yes | Yes | Fast | Active |
| ge-0/0/2 | Actor | No | No | Yes | Yes | Yes | Yes | Fast | Active |
| ge-0/0/2 | Partner | No | No | Yes | Yes | Yes | Yes | Fast | Active |
| ge-0/0/3 | Actor | No | No | Yes | Yes | Yes | Yes | Fast | Active |
| ge-0/0/3 | Partner | No | No | Yes | Yes | Yes | Yes | Fast | Active |
| LACP protocol: | Receive State | | | Transmit State | | | Mux State | | |
| ge-0/0/1 | Current | | | Fast periodic | | | Collecting distributing | | |
| ge-0/0/2 | Current | | | Fast periodic | | | Collecting distributing | | |
| ge-0/0/3 | Current | | | Fast periodic | | | Collecting distributing | | |
| LACP info: | Role | System | | | System | | | Port | Port |
| Port | | | | | | | | | |
| | | | priority | | identifier | | priority | | number |
| key | | | | | | | | | |
| 1 | ge-0/0/1 | Actor | 127 | 00:05:86:4e:b6:c0 | | | 127 | 1 | |
| 1 | ge-0/0/1 | Partner | 127 | 00:05:86:7e:d3:c0 | | | 127 | 1 | |
| 1 | ge-0/0/2 | Actor | 127 | 00:05:86:4e:b6:c0 | | | 127 | 2 | |
| 1 | ge-0/0/2 | Partner | 127 | 00:05:86:7e:d3:c0 | | | 127 | 2 | |
| 1 | ge-0/0/3 | Actor | 127 | 00:05:86:4e:b6:c0 | | | 127 | 3 | |
| 1 | ge-0/0/3 | Partner | 127 | 00:05:86:7e:d3:c0 | | | 127 | 3 | |
| 1 | | | | | | | | | |

show lacp interfaces (Gigabit Ethernet)

```
user@host> show lacp interfaces ge-0/3/0
```

```
Aggregated interface: ae0
LACP State:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
ge-0/3/0        Actor  No   No   Yes  Yes  Yes  Yes    Fast    Active
ge-0/3/0        Partner No   No   Yes  Yes  Yes  Yes    Fast    Active
LACP Protocol:  Receive State  Transmit State      Mux State
ge-0/3/0        Current  Fast periodic  Collecting distributing
```

show lacp interfaces (10 Gigabit Ethernet)

```
user@host> show lacp interfaces xe-1/0/2
```

```
Aggregated interface: ae0
LACP State:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
xe-1/0/2        Actor  No   No   Yes  Yes  Yes  Yes    Fast    Active
xe-1/0/2        Partner No   No   Yes  Yes  Yes  Yes    Fast    Active
LACP Protocol:  Receive State  Transmit State      Mux State
xe-1/0/2        Current  Fast periodic  Collecting distributing
```

show lacp statistics interfaces (View)

| | |
|---------------------------------|---|
| Syntax | show lacp statistics interfaces <i>interface-name</i> |
| Release Information | <p>Command modified in Release 10.2 of Junos OS.</p> <p>Command introduced in Release 11.1 of Junos OS for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display Link Aggregation Control Protocol (LACP) statistics about the specified aggregated Ethernet interface or redundant Ethernet interface. If you do not specify an interface name, LACP statistics for all interfaces are displayed. |
| Options | <i>interface-name</i> —(Optional) Name of an interface. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Verifying LACP on Redundant Ethernet Interfaces • Verifying the Status of a LAG Interface on page 90 • Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 91 • Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch • Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch |
| List of Sample Output | show lacp statistics interfaces on page 340 |
| Output Fields | <p>Table 34 on page 339 lists the output fields for the show lacp statistics interfaces command. Output fields are listed in the approximate order in which they appear.</p> |

Table 34: show lacp statistics interfaces Output Fields

| Field Name | Field Description |
|----------------------|-----------------------------|
| Aggregated interface | Aggregated interface value. |

Table 34: show lacp statistics interfaces Output Fields (continued)

| Field Name | Field Description |
|------------------------|--|
| LACP Statistics | <p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> • LACP Rx—LACP received counter that increments for each normal hello. • LACP Tx—Number of LACP transmit packet errors logged. • Unknown Rx—Number of unrecognized packet errors logged. • Illegal Rx—Number of invalid packets received. <p>NOTE: Starting in Junos OS Evolved Release 18.3R1, the clear interfaces statistics command clears LACP statistics as well as the counters displayed in the show lacp statistics interfaces command.</p> |

Sample Output

show lacp statistics interfaces

```
user@host> show lacp statistics interfaces ae0
```

```
Aggregated interface: ae0
```

| LACP Statistics: | LACP Rx | LACP Tx | Unknown Rx | Illegal Rx |
|------------------|---------|---------|------------|------------|
| ge-2/0/0 | 1352 | 2035 | 0 | 0 |
| ge-2/0/1 | 1352 | 2056 | 0 | 0 |
| ge-2/2/0 | 1352 | 2045 | 0 | 0 |
| ge-2/2/1 | 1352 | 2043 | 0 | 0 |

CHAPTER 15

Redundant Trunk Group Operational Command

- `show redundant-trunk-group`

show redundant-trunk-group

| | |
|---------------------------------|---|
| Syntax | show redundant-trunk-group <group-name group-name> |
| Release Information | Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X50-D15 for the QFX Series. |
| Description | Display information about redundant trunk groups. |
| Options | group-name group-name —Display information about the specified redundant trunk group. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • <i>Example: Configuring Redundant Trunk Links for Faster Recovery on EX Series Switches</i> • <i>Example: Configuring Redundant Trunk Links for Faster Recovery on Devices with ELS Support</i> • <i>Understanding Redundant Trunk Links (Legacy RTG Configuration)</i> |
| List of Sample Output | show redundant-trunk-group group-name Group1 on page 343 |
| Output Fields | Table 35 on page 342 lists the output fields for the show redundant-trunk-group command. Output fields are listed in the approximate order in which they appear. |

Table 35: show redundant-trunk-group Output Fields

| Field Name | Field Description |
|--------------------------|--|
| Group name | Name of the redundant trunk port group. |
| Interface | Name of an interface belonging to the trunk port group. |
| State | Operating state of the interface. <ul style="list-style-type: none"> • Up denotes the interface is up. • Down denotes the interface is down. • Pri denotes a primary interface. • Act denotes an active interface. |
| Time of last flap | Date and time at which the advertised link became unavailable, and then, available again. |
| Flap count | Total number of flaps since the last switch reboot. |

Sample Output

`show redundant-trunk-group group-name Group1`

```
user@switch> show redundant-trunk-group group-name Group1
```

| Group name | Interface | State | Time of last flap | Flap Count |
|------------|-------------|------------|-------------------|------------|
| Group1 | ge-0/0/45.0 | UP/Pri/Act | Never | 0 |
| | ge-0/0/47.0 | UP | Never | 0 |

