

Interfaces Feature Guide for the QFX Series

Release
15.1



Modified: 2017-01-10

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

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

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

Interfaces Feature Guide for the QFX Series

15.1

Copyright © 2016, Juniper Networks, Inc.

All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

| | | |
|------------------|---|----------|
| | About the Documentation | xvii |
| | Documentation and Release Notes | xvii |
| | Supported Platforms | xvii |
| | Using the Examples in This Manual | xvii |
| | Merging a Full Example | xviii |
| | Merging a Snippet | xviii |
| | Documentation Conventions | xix |
| | Documentation Feedback | xxi |
| | Requesting Technical Support | xxi |
| | Self-Help Online Tools and Resources | xxi |
| | Opening a Case with JTAC | xxii |
| Part 1 | Interfaces | |
| Chapter 1 | Understanding Interfaces | 3 |
| | Interfaces Overview | 3 |
| | Network Interfaces | 3 |
| | Special Interfaces | 4 |
| | Understanding Interface Naming Conventions | 5 |
| | Physical Part of an Interface Name | 6 |
| | Logical Part of an Interface Name on a Switch Running QFabric Software Package | 15 |
| | Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software | 15 |
| | Wildcard Characters in Interface Names | 16 |
| | Understanding Interface Ranges | 16 |
| | Understanding Management Interfaces | 17 |
| | Understanding Port Ranges and System Modes | 18 |
| | Port Ranges for Different Media Types | 19 |
| | Supported System Modes | 42 |
| | Configuring the Port Type on QFX3600 Standalone Switches | 44 |
| | Configuring the QSFP+ Port Type on QFX3500 Standalone Switches | 46 |
| | Configuring the QSFP+ Port Type on QFX5100 Devices | 48 |
| | Configuring the Interface Address | 51 |
| | Configuring Gigabit and 10-Gigabit Ethernet Interfaces | 53 |
| | Configuring Port Mode on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 54 |
| | Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches | 55 |
| | Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches | 55 |

| | | |
|------------------|--|-----------|
| | Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches | 57 |
| | Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches | 57 |
| | Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches | 58 |
| | Configuring Ethernet Loopback Capability | 58 |
| | Configuring an LPM Table With Junos OS Release 13.2X51-D10 | 60 |
| | Channelizing Interfaces | 63 |
| | Channelizing Interfaces on QFX5110-48S Switches | 73 |
| | Channelizing Interfaces on QFX5200 Switches | 75 |
| | Configuring the System Mode | 77 |
| | Monitoring Interface Status and Traffic | 79 |
| | Troubleshooting Network Interfaces | 80 |
| | The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP or SFP+ module is down | 80 |
| Part 2 | Ethernet OAM Link Fault Management | |
| Chapter 2 | Understanding Ethernet OAM Link Fault Management | 83 |
| | Understanding Ethernet OAM Link Fault Management for an EX Series Switch | 83 |
| | Configuring Ethernet OAM Link Fault Management (CLI Procedure) | 84 |
| | Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches | 86 |
| Chapter 3 | Ethernet OAM Link Fault Management Operational Commands | 91 |
| | show oam ethernet link-fault-management | 92 |
| Part 3 | Generic Routing Encapsulation (GRE) | |
| Chapter 4 | Understanding GRE | 99 |
| | Understanding Generic Routing Encapsulation | 99 |
| | Overview of GRE | 99 |
| | GRE Tunneling | 100 |
| | Encapsulation and De-Encapsulation on the Switch | 100 |
| | Number of Source and Destination Tunnels Allowed on a Switch | 100 |
| | Class of Service on GRE Tunnels | 101 |
| | Applying Firewall Filters to GRE Traffic | 101 |
| | Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches | 101 |
| | Configuration Limitations | 102 |
| | Configuring Generic Routing Encapsulation Tunneling | 102 |
| | Configuring a GRE Tunnel | 103 |
| | Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly | 103 |

| | | |
|------------------|--|------------|
| Part 4 | IP Directed Broadcast | |
| Chapter 5 | Understanding IP Directed Broadcast | 107 |
| | Understanding IP Directed Broadcast | 107 |
| | IP Directed Broadcast Overview | 107 |
| | IP Directed Broadcast Implementation | 108 |
| | When to Enable IP Directed Broadcast | 108 |
| | When Not to Enable IP Directed Broadcast | 108 |
| | Configuring IP Directed Broadcast (CLI Procedure) | 109 |
| | Example: Configuring IP Directed Broadcast | 110 |
| Part 5 | Layer 3 Logical Interfaces | |
| Chapter 6 | Understanding Layer 3 Logical Interfaces | 117 |
| | Understanding Layer 3 Logical Interfaces | 117 |
| | Configuring a Layer 3 Logical Interface | 118 |
| | Verifying That Layer 3 Logical Interfaces Are Working | 118 |
| Part 6 | Local Link Bias | |
| Chapter 7 | Understanding Local Link Bias | 123 |
| | Understanding Local Link Bias | 123 |
| | Configuring Local Link Bias (CLI Procedure) | 125 |
| Part 7 | Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP) | |
| Chapter 8 | Understanding LAGs and LACP | 129 |
| | Understanding Aggregated Ethernet Interfaces and LACP | 129 |
| | Link Aggregation Group | 130 |
| | Link Aggregation Control Protocol (LACP) | 131 |
| | Configuring Aggregated Ethernet LACP | 132 |
| | Configuring Link Aggregation | 133 |
| | Creating an Aggregated Ethernet Interface | 134 |
| | Configuring the VLAN Name and VLAN ID Number | 134 |
| | Configuring Aggregated Ethernet LACP | 135 |
| | Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch | 136 |
| | Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch | 140 |
| | Verifying the Status of a LAG Interface | 145 |
| | Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets | 145 |
| | Verifying the LACP Setup | 145 |
| | Verifying That LACP Packets Are Being Exchanged | 146 |
| | Troubleshooting an Aggregated Ethernet Interface | 146 |

| | | |
|-------------------|---|------------|
| Part 8 | Redundant Trunk Groups | |
| Chapter 9 | Understanding Redundant Trunk Groups | 151 |
| | Understanding Redundant Trunk Links | 151 |
| | Example: Configuring Redundant Trunk Links for Faster Recovery | 153 |
| Part 9 | Resilient Hashing | |
| Chapter 10 | Understanding Resilient Hashing | 161 |
| | Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop | |
| | ECMP Traffic | 161 |
| | Understanding the Hashing Algorithm | 161 |
| | IP (IPv4 and IPv6) | 162 |
| | MPLS | 164 |
| | MAC-in-MAC Packet Hashing | 165 |
| | Layer 2 Header Hashing | 166 |
| | Understanding the Use of Resilient Hashing to Minimize Flow Remapping in | |
| | Trunk/ECMP Groups | 167 |
| | Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing | 167 |
| | Limitations and Caveats for Resilient Hashing | 168 |
| | Resilient Hashing on LAGs | 168 |
| | Resilient Hashing on ECMP | 169 |
| | Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP | |
| | Traffic (CLI Procedure) | 169 |
| | Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing | 169 |
| | Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing | 170 |
| | Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing | 170 |
| | Configuring Resilient Hashing for Trunk/ECMP Groups | 171 |
| | Configuring Resilient Hashing on LAGs | 171 |
| | Configuring Resilient Hashing on ECMP Groups | 171 |
| Part 10 | Uplink Failure Detection | |
| Chapter 11 | Understanding Uplink Failure Detection | 175 |
| | Overview of Uplink Failure Detection | 175 |
| | Uplink Failure Detection Configuration | 175 |
| | Failure Detection Pair | 176 |
| | Configuring Interfaces for Uplink Failure Detection | 177 |
| | Example: Configuring Interfaces for Uplink Failure Detection | 178 |
| | Verifying That Uplink Failure Detection Is Working Correctly | 182 |
| Part 11 | Configuration Statements and Operational Commands | |
| Chapter 12 | Ethernet OAM Link Fault Management Configuration Statements | 185 |
| | action (OAM LFM) | 186 |
| | action-profile | 187 |

| | | |
|-------------------|--|------------|
| | allow-remote-loopback | 188 |
| | ethernet (OAM LFM) | 189 |
| | event-thresholds | 191 |
| | event (OAM LFM) | 192 |
| | frame-error | 192 |
| | frame-period | 193 |
| | frame-period-summary | 193 |
| | oam | 194 |
| | interface (OAM LFM) | 196 |
| | link-adjacency-loss | 197 |
| | link-discovery | 197 |
| | link-down | 198 |
| | link-event-rate | 198 |
| | link-fault-management | 199 |
| | negotiation-options | 200 |
| | no-allow-link-events | 200 |
| | pdu-interval | 201 |
| | pdu-threshold | 201 |
| | remote-loopback | 202 |
| | symbol-period | 202 |
| | syslog (OAM LFM) | 203 |
| | traceoptions (OAM LFM) | 204 |
| Chapter 13 | GRE Configuration Statements | 207 |
| | destination (Tunnels) | 207 |
| | tunnel | 208 |
| | tunnel-port | 208 |
| Chapter 14 | Interfaces Configuration Statements | 209 |
| | address | 211 |
| | alarm (chassis) | 215 |
| | auto-negotiation | 216 |
| | ccc | 217 |
| | channel-speed | 218 |
| | configured-flow-control | 219 |
| | craft-lockout | 220 |
| | description | 221 |
| | ethernet (Alarm) | 222 |
| | ethernet-switching | 223 |
| | ether-options | 224 |
| | eui-64 | 225 |
| | family | 226 |
| | fibre-channel (Alarm) | 230 |
| | filter | 231 |
| | flow-control | 233 |
| | fpc | 234 |
| | gratuitous-arp-reply | 235 |
| | hold-time (Physical Interface) | 236 |
| | irb (Interfaces) | 238 |
| | inet (interfaces) | 241 |

| | | |
|-------------------|---|------------|
| | inet6 (interfaces) | 242 |
| | interface-mode | 243 |
| | interface-range | 245 |
| | interfaces | 247 |
| | link-down | 254 |
| | link-mode | 255 |
| | link-speed | 256 |
| | loopback (Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet) | 257 |
| | mac | 257 |
| | management-ethernet (Alarm) | 258 |
| | member | 259 |
| | member-range | 260 |
| | mtu | 261 |
| | no-gratuitous-arp-request | 262 |
| | pic | 263 |
| | rx-buffers | 264 |
| | source | 265 |
| | speed | 266 |
| | targeted-broadcast | 267 |
| | traceoptions (Individual Interfaces) | 268 |
| | tx-buffers | 269 |
| | unit | 271 |
| | vlan-id | 272 |
| | vlan-tagging | 273 |
| Chapter 15 | IP Directed Broadcast Configuration Statement | 275 |
| | targeted-broadcast | 275 |
| Chapter 16 | LAGs and LACP Configuration Statements | 277 |
| | aggregated-devices | 278 |
| | aggregated-ether-options | 279 |
| | chassis | 281 |
| | 802.3ad | 282 |
| | device-count | 283 |
| | ethernet | 283 |
| | force-up | 284 |
| | lACP (802.3ad) | 285 |
| | lACP (Aggregated Ethernet) | 286 |
| | periodic | 287 |
| Chapter 17 | Local Link Bias Configuration Statements | 289 |
| | local-bias | 289 |
| Chapter 18 | Redundant Trunk Groups Configuration Statements | 291 |
| | group (Redundant Trunk Groups) | 292 |
| | interface (Redundant Trunk Groups) | 293 |
| | preempt-cutover-timer | 294 |
| | redundant-trunk-group | 295 |

| | | |
|-------------------|--|------------|
| Chapter 19 | Resilient Hashing Configuration Statements | 297 |
| | ecmp-resilient-hash | 297 |
| | enhanced-hash-key | 298 |
| | hash-mode | 301 |
| | inet (enhanced-hash-key) | 303 |
| | inet6 (enhanced-hash-key) | 305 |
| | resilient-hash | 307 |
| Chapter 20 | Uplink Failure Detection Configuration Statements | 309 |
| | group | 309 |
| | link-to-disable | 310 |
| | link-to-monitor | 310 |
| | uplink-failure-detection | 311 |
| Chapter 21 | Interfaces Operational Commands | 313 |
| | monitor interface | 314 |
| | show interfaces diagnostics optics | 326 |
| | show interfaces ge | 340 |
| | show interfaces (GRE) | 355 |
| | show interfaces irb | 364 |
| | show interfaces queue | 370 |
| | show interfaces xe | 409 |
| Chapter 22 | LAGs and LACP Operational Commands | 427 |
| | show lacp interfaces | 428 |
| | show lacp statistics interfaces (View) | 433 |
| Chapter 23 | Redundant Trunk Group Operational Commands | 435 |
| | show redundant-trunk-group | 436 |
| Chapter 24 | Resilient Hashing Operational Commands | 439 |
| | show forwarding-options enhanced-hash-key | 440 |
| Chapter 25 | Uplink Failure Detection Operational Commands | 445 |
| | show uplink-failure-detection | 446 |

List of Figures

| | | |
|-------------------|--|------------|
| Part 4 | IP Directed Broadcast | |
| Chapter 5 | Understanding IP Directed Broadcast | 107 |
| | Figure 1: Topology for IP Directed Broadcast | 111 |
| Part 6 | Local Link Bias | |
| Chapter 7 | Understanding Local Link Bias | 123 |
| | Figure 2: Egress Traffic Flow with Local Link Bias | 123 |
| | Figure 3: Egress Traffic Flow without Local Link Bias | 124 |
| Part 8 | Redundant Trunk Groups | |
| Chapter 9 | Understanding Redundant Trunk Groups | 151 |
| | Figure 4: Redundant Trunk Group, Link 1 Active | 152 |
| | Figure 5: Redundant Trunk Group, Link 2 Active | 152 |
| | Figure 6: Topology for Configuring the Redundant Trunk Links | 155 |
| Part 10 | Uplink Failure Detection | |
| Chapter 11 | Understanding Uplink Failure Detection | 175 |
| | Figure 7: Uplink Failure Detection Configuration on Switches | 176 |
| | Figure 8: Uplink Failure Detection Configuration on Switches | 179 |

List of Tables

| | | |
|------------------|---|-------------|
| | About the Documentation | xvii |
| | Table 1: Notice Icons | xix |
| | Table 2: Text and Syntax Conventions | xix |
| Part 1 | Interfaces | |
| Chapter 1 | Understanding Interfaces | 3 |
| | Table 3: Network Interface Types and Purposes | 4 |
| | Table 4: Special Interface Types and Purposes | 5 |
| | Table 5: Valid Port Ranges on QFX3500 Switches Running QFabric Software Package | 20 |
| | Table 6: Valid Port Ranges on QFX3500 Switches Running Enhanced Layer 2 Software | 25 |
| | Table 7: Valid Port Ranges on QFX3600 Switches Running QFabric Software Package | 28 |
| | Table 8: Valid Port Ranges on QFX3600 Switches Running Enhanced Layer 2 Software | 31 |
| | Table 9: Valid Port Ranges on QFX3600 Node Devices Running QFabric Software Package | 34 |
| | Table 10: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running Enhanced Layer 2 Software | 36 |
| | Table 11: Valid Port Ranges on QFX5100-48S and QFX5100-48T Switches Running QFabric Software Package | 39 |
| | Table 12: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software | 43 |
| | Table 13: Example LPM Table Combinations Using l2-profile-one With Junos OS 13.2X51-D10 | 61 |
| | Table 14: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10 | 62 |
| | Table 15: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings | 65 |
| | Table 16: QFX10002-72Q Switch Port Mappings | 67 |
| | Table 17: System Modes Supported on QFX5100 Switches with QFX-EM-4Q or QFX-PFA-4Q Expansion Modules Installed | 77 |
| | Table 18: System Modes Supported on QFX5100-24Q Switches with the EX4600-8F Expansion Module Installed | 78 |
| | Table 19: System Modes Supported on QFX5100-24Q Switches with EX4600-8F and QFX-EM-4Q Expansion Modules Installed | 79 |

| | | |
|-------------------|--|------------|
| Part 2 | Ethernet OAM Link Fault Management | |
| Chapter 3 | Ethernet OAM Link Fault Management Operational Commands | 91 |
| | Table 20: show oam ethernet link-fault-management Output Fields | 92 |
| Part 3 | Generic Routing Encapsulation (GRE) | |
| Chapter 4 | Understanding GRE | 99 |
| | Table 21: Firewall Filter Application Points for Tunneled Packets | 101 |
| | Table 22: Features Not Supported with GRE | 102 |
| Part 4 | IP Directed Broadcast | |
| Chapter 5 | Understanding IP Directed Broadcast | 107 |
| | Table 23: Components of the IP Directed Broadcast Topology | 111 |
| Part 7 | Link Aggregation Groups (LAGs) and Link Aggregation Control Protocol (LACP) | |
| Chapter 8 | Understanding LAGs and LACP | 129 |
| | Table 24: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch | 137 |
| Part 8 | Redundant Trunk Groups | |
| Chapter 9 | Understanding Redundant Trunk Groups | 151 |
| | Table 25: Components of the Redundant Trunk Link Topology | 156 |
| Part 9 | Resilient Hashing | |
| Chapter 10 | Understanding Resilient Hashing | 161 |
| | Table 26: IPv4 and IPv6 Hashing Fields | 163 |
| | Table 27: MPLS Hashing Fields | 164 |
| | Table 28: MAC-in-MAC Hashing Fields | 165 |
| | Table 29: Layer 2 Header Hashing Fields | 166 |
| | Table 30: Destination Path Results for Static Hashing and for Resilient Hashing When Members Are Added to or Deleted from Trunk Groups | 168 |
| Part 10 | Uplink Failure Detection | |
| Chapter 11 | Understanding Uplink Failure Detection | 175 |
| | Table 31: Settings for Uplink Failure Protection Example | 179 |
| Part 11 | Configuration Statements and Operational Commands | |
| Chapter 21 | Interfaces Operational Commands | 313 |
| | Table 32: Output Control Keys for the monitor interface interface-name Command | 314 |
| | Table 33: Output Control Keys for the monitor interface traffic Command | 315 |
| | Table 34: monitor interface Output Fields | 316 |
| | Table 35: show interfaces diagnostics optics Output Fields | 326 |

| | | |
|-------------------|---|------------|
| | Table 36: show interfaces ge Output Fields | 341 |
| | Table 37: GRE show interfaces Output Fields | 356 |
| | Table 38: show interfaces irb Output Fields | 364 |
| | Table 39: Layer 2 Overhead and Transmitted Packets or Byte Counts | 371 |
| | Table 40: show interfaces queue Output Fields | 374 |
| | Table 41: Byte Count by Interface Hardware | 377 |
| | Table 42: show interfaces xe Output Fields | 410 |
| Chapter 22 | LAGs and LACP Operational Commands | 427 |
| | Table 43: show lacp interfaces Output Fields | 429 |
| | Table 44: show lacp statistics interfaces Output Fields | 433 |
| Chapter 23 | Redundant Trunk Group Operational Commands | 435 |
| | Table 45: show redundant-trunk-group Output Fields | 436 |
| Chapter 24 | Resilient Hashing Operational Commands | 439 |
| | Table 46: show forwarding-options enhanced-hash-key Output Fields | 440 |
| Chapter 25 | Uplink Failure Detection Operational Commands | 445 |
| | Table 47: show uplink-failure-detection Output Fields | 446 |

About the Documentation

- Documentation and Release Notes on page xvii
- Supported Platforms on page xvii
- Using the Examples in This Manual on page xvii
- Documentation Conventions on page xix
- Documentation Feedback on page xxi
- Requesting Technical Support on page xxi

Documentation and Release Notes

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

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

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

Supported Platforms

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

- [QFX Series](#)

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 xix](#) defines notice icons used in this guide.

Table 1: Notice Icons

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

[Table 2 on page xix](#) defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

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

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

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

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

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

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

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

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

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

Opening a Case with JTAC

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

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

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

PART 1

Interfaces

- [Understanding Interfaces on page 3](#)

CHAPTER 1

Understanding Interfaces

- [Interfaces Overview on page 3](#)
- [Understanding Interface Naming Conventions on page 5](#)
- [Understanding Interface Ranges on page 16](#)
- [Understanding Management Interfaces on page 17](#)
- [Understanding Port Ranges and System Modes on page 18](#)
- [Configuring the Port Type on QFX3600 Standalone Switches on page 44](#)
- [Configuring the QSFP+ Port Type on QFX3500 Standalone Switches on page 46](#)
- [Configuring the QSFP+ Port Type on QFX5100 Devices on page 48](#)
- [Configuring the Interface Address on page 51](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53](#)
- [Configuring Ethernet Loopback Capability on page 58](#)
- [Configuring an LPM Table With Junos OS Release 13.2X51-D10 on page 60](#)
- [Channelizing Interfaces on page 63](#)
- [Channelizing Interfaces on QFX5110-48S Switches on page 73](#)
- [Channelizing Interfaces on QFX5200 Switches on page 75](#)
- [Configuring the System Mode on page 77](#)
- [Monitoring Interface Status and Traffic on page 79](#)
- [Troubleshooting Network Interfaces on page 80](#)

Interfaces Overview

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

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

Network Interfaces

Network interfaces connect to the network and carry network traffic. [Table 3 on page 4](#) lists the types of network interfaces supported.

Table 3: Network Interface Types and Purposes

| Type | Purpose |
|---|---|
| Aggregated Ethernet interfaces | You can group Ethernet interfaces at the physical layer to form a single link-layer interface, also known as a <i>link aggregation group (LAG)</i> or <i>bundle</i> . These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth. |
| 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 | You can 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 | You can 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. |
| LAN access interfaces | You can 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 | You can 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 | You can 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 | You can 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. |
| 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, 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. |

Special Interfaces

Table 4 on page 5 lists the types of special interfaces supported.

Table 4: Special Interface Types and Purposes

| Type | Purpose |
|---|---|
| Console port | Each device has a serial console port, labeled 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> |

Related Documentation

- [Understanding Aggregated Ethernet Interfaces and LACP on page 129](#)
- [Understanding Interface Naming Conventions on page 5](#)
- [Understanding Layer 3 Logical Interfaces on page 117](#)
- [Understanding Management Interfaces on page 17](#)
- *Understanding Integrated Routing and Bridging*
- *Overview of Fibre Channel*

Understanding Interface Naming Conventions

The QFX Series and the EX4600 device uses a naming convention for defining the interfaces that is similar to that of other platforms running under Juniper Networks Junos OS. This topic provides brief information about the naming conventions used for interfaces on the QFX Series and on EX4600 switches.

This topic describes:

- [Physical Part of an Interface Name on page 6](#)
- [Logical Part of an Interface Name on a Switch Running QFabric Software Package on page 15](#)
- [Logical Part of a Channelized Interface Name on a Switch Running Enhanced Layer 2 Software on page 15](#)
- [Wildcard Characters in Interface Names on page 16](#)

Physical Part of an Interface Name

Interfaces in Junos OS are specified as follows:

device-name:type-fpc/pic/port

The convention is as follows:

- *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
 - **xle**—40-Gigabit Ethernet interface (QFX3500, QFX3600, and QFX5100 switches running a QFabric software package)
 - **et**—40-Gigabit Ethernet interface (QFX3500, QFX3600, QFX5100, 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:
 - On a QFX3500 switch running a 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.
 - On a QFX3500 switch running 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.
 - On a QFX3500 Node device running a QFabric software package, PIC 0 can support 48 ports and PIC 1 can support four 40-Gigabit data plane uplink ports.
 - On a QFX3600 switch running a QFabric software package, PIC 0 can support 64 10-Gigabit Ethernet ports, and PIC 1 can support 16 40-Gigabit Ethernet ports.
 - On a QFX3600 switch running Enhanced Layer 2 software, PIC 0 can support 64 10-Gigabit Ethernet ports and can also support 16 40-Gigabit Ethernet ports.
 - On a 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.
 - On a QFX5100-48S switch running Enhanced Layer 2 software, PIC 0 provides six 40-Gbps QSFP+ ports and 48 10-Gigabit Ethernet interfaces.
 - On an EX4600 device running 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 provides 4 40-Gbps QSFP+ ports. The EX4600-EM-8F expansion module provides 8 40-Gbps QSFP+ 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.
 - On a QFX5100-48S switch running a QFabric software package, PIC 1 provides six 40-Gbps QSFP+ ports, and PIC 0 provides 48 10-Gigabit Ethernet interfaces.
 - On a QFX5100-24Q switch running 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.
 - On a QFX5100-96S switch running Enhanced Layer 2 software, PIC 0 provides 96 10-Gigabit Ethernet interfaces and 8 40-Gbps QSFP+ ports.
 - On a QFX5110-48S switch running 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.

- On a QFX5200-32C switch running 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.
- On a QFX10002-36Q switch running Enhanced Layer 2 software, PIC 0 provides 144 10-Gigabit Ethernet interfaces, and 36 40-Gbps QSFP+ ports, and 12 100-Gigabit Ethernet interfaces.
- On a QFX10002-72Q switch running Enhanced Layer 2 software, PIC 0 provides 288 10-Gigabit Ethernet interfaces, and 72 40-Gbps QSFP+ ports, and 24 100-Gigabit Ethernet interfaces.
- On a QFX10008 switch running 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.
- On a QFX10016 switch running 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:
 - On a QFX3500 switch running a QFabric software package, 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.

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 [“Configuring the QSFP+ Port Type on QFX3500 Standalone Switches” on page 46](#)).

- On a QFX3500 switch running 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 [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On a QFX3600 switch running a QFabric software package, 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.

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 [“Configuring the Port Type on QFX3600 Standalone Switches” on page 44](#)) and use QSFP+ to four SFP+ copper breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches.

- On a QFX3600 Node device running a QFabric software package, 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.

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 *Configuring the Port Type on QFX3600 Node Devices*).

- On a QFX3600 switch running 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 [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On a QFX5100-48S switch running 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 [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On an EX4600 switch running 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 40-Gbps QSFP+ 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 [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.
- On a QFX5100-48S switch running 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 *Configuring the QSFP+ Port Type on QFX5100 Switches* for information on how to configure the port mode of 40-Gbps QSFP+ ports.
- On a QFX5100-24Q switch running Enhanced Layer 2 software, 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 [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.



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.

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 [“Configuring the System Mode” on page 77](#) for information on how to configure the system mode.

- On a QFX5100-96S switch running Enhanced Layer 2 software, PIC 0 can support 96 10-Gigabit Ethernet ports labeled 0 through 95, and 8 40-Gbps QSFP+ ports labeled 96 through 103. See [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.



NOTE: You can only channelize the 40-Gbps QSFP+ ports provided in ports 96 and 100, because only 104 logical ports can be channelized.

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 [“Configuring the System Mode” on page 77](#) for information on how to configure the system mode.

- On a QFX5110-48S switch running 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. 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.
- On a QFX5200-32C switch running Enhanced Layer 2 software, 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.

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 [“Channelizing Interfaces on QFX5200 Switches” on page 75](#) for information on how to configure and channelize the interfaces.



NOTE: Autochannelization is not supported.

- On a QFX10002-36Q switch running Enhanced Layer 2 software, 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.

Each QSFP28 socket can be configured to support:

- 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.

Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

Each of the 12 QSFP28 ports support:

- 100-Gigabit Ethernet QSFP28 transceivers
- 40-Gigabit Ethernet QSFP+ transceivers

Each of the 36 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers
- Access ports
- On a QFX10002-72Q switch running Enhanced Layer 2 software, 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.

Each QSFP28 socket can be configured to support:

- 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.

Any of the 72 ports 0 through 71 can be configured as either uplink or access ports. See [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

Each of the 24 QSFP28 ports support:

- 100-Gigabit Ethernet QSFP28 transceivers

Each of the 72 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers

Each of the 36 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers
- Access ports
- Uplink ports
- On a QFX10008 switch running 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
 - QFX10000-30C, a 30-port 100-Gigabit or 40-Gigabit Ethernet QSFP28 line card

The QFX10000-36Q line cards supports

- 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.
- Each QSFP28 socket can be configured to support:
 - 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.

Any of the 36 ports 0 through 35 can be configured as either uplink or access ports. See [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

- Each of the 12 QSFP28 ports supports:
 - 100-Gigabit Ethernet QSFP28 transceivers

- 40-Gigabit Ethernet QSFP+ transceivers
- Each of the 12 QSFP28 ports supports:
 - 100-Gigabit Ethernet QSFP28 transceivers
 - 40-Gigabit Ethernet QSFP+ transceivers

Each of the 36 QSFP+ ports support:

- 40-Gigabit Ethernet QSFP+ transceivers
- Access ports
- Uplink ports

The QFX10000-30C line cards supports

- Thirty 28-Gbps QSFP+ Pluggable Solution (QSFP28) cages that support either 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The QFX10000-30C ports auto detect the type of transceiver installed and set the configuration to the appropriate speed.
- Each QSFP28 socket can be configured to support:
 - 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.

See [“Channelizing Interfaces” on page 63](#) for information on how to configure and channelize the 40-Gbps QSFP+ ports.

- Each of the 30 QSFP28 ports supports:
 - 100-Gigabit Ethernet QSFP28 transceivers
 - 40-Gigabit Ethernet QSFP+ transceivers
 - Access ports
 - Uplink ports
- On a QFX10016 switch running Enhanced Layer 2 software, there are 16 slots, which you can populate with two types line cards:
 - QFX10000-36Q, a 36-port 40-Gigabit Ethernet quad small form-factor pluggable plus transceiver (QSFP+) or 12-port 100GbE QSFP28 line card

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.

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.

You can use 40-Gigabit Ethernet using QSFP+ optical transceivers.

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.

Any of the 36 ports 0 through 35 can be configured as either uplink or access ports.

Each of the 12 QSFP28 ports supports:

- 100-Gigabit Ethernet QSFP28 transceivers
- 40-Gigabit Ethernet QSFP+ transceivers

Each of the 36 QSFP+ ports supports:

- 40-Gigabit Ethernet QSFP+ transceivers
- Access ports

You can use 40-Gigabit Ethernet QSFP+ transceivers in any downstream port.

- Uplink ports

You can configure all the QSFP+ ports as uplinks.

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.

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 `set chassis fpc 0 pic 0 port 1 channel-speed 10g` command. For example, to channelize the first switch port, use the `set chassis fpc 0 pic 0 port 1 channel-speed 10g` command.

- The QFX10000-30C line card consists of thirty 28-Gbps QSFP+ Pluggable Solution (QSFP28) cages that support either 40-Gigabit Ethernet or 100-Gigabit Ethernet optical transceivers. The QFX10000-30C ports auto

detect the type of transceiver installed and set the configuration to the appropriate speed.

Each QSFP28 socket supports:

- 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.

Any of the 30 ports 0 through 29 can be configured as either uplink or access ports, and of the 30 QSFP28 ports supports:

- 100-Gigabit Ethernet QSFP28 transceivers
- 40-Gigabit Ethernet QSFP+ transceivers

Logical Part of an Interface Name on a Switch Running QFabric Software Package

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

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

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

Related Documentation

- [Interfaces Overview on page 3](#)
- [Channelizing Interfaces on page 63](#)
- [Configuring the System Mode on page 77](#)
- [Understanding Management Interfaces on page 17](#)
- [Understanding Port Ranges and System Modes on page 18](#)
- *Rear Panel of a QFX3500 Device*
- *Front Panel of a QFX3600 Device*
- *Junos OS Network Interfaces Library for Routing Devices*

Understanding Interface Ranges

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

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



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

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

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



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

- **protocols isis interface**

- [protocols sflow interfaces](#)

Related Documentation

- [Interfaces Overview on page 3](#)
- [Interfaces Overview](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [Configuring Link Aggregation on page 133](#)
- [Configuring a Layer 3 Logical Interface on page 118](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)
- [interface-range on page 245](#)

Understanding Management Interfaces

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



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

Management interface port ranges vary based on device type:

- QFX3500 devices:

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

- QFX3600 devices:

There are two RJ-45 management interfaces (labeled **C0** and **C1**) and two SFP management interfaces (labeled **C0S** and **C1S**). On a QFX3600 standalone switch, the RJ-45 management interfaces and SFP management interfaces correspond to **me0** and **me1**. On a QFX3600 Node device, the RJ-45 management interfaces and SFP management interfaces correspond to **me5** and **me6**. Each pair of management interfaces correspond to one Ethernet interface—for example, both RJ-45 management interfaces (labeled **C0** and **C0s**) can correspond to **me0**, and both SFP management

interfaces (labeled **C1** and **C1S**) can correspond to me1. By default, both RJ-45 management interfaces are active. If you insert an SFP interface into the SFP management port (**C0S**, for example), the SFP interface would become the active management interface, and the corresponding RJ-45 management interface (**C0**) is disabled.



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

- On QFX5100, 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.

**Related
Documentation**

- [Interfaces Overview on page 3](#)

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 19](#)
- [Supported System Modes on page 42](#)

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 7 on page 28](#) 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 8 on page 31](#) 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 10 on page 36](#) 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 10 on page 36](#) 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 11 on page 39](#) for physical port to logical port mappings.

- For QFX5100-24Q and QFX5100-96S switches running Enhanced Layer 2 Software, see [Table 12 on page 43](#) for physical port to logical port mappings for different system modes.

Table 5: 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 5: 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 5: 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 5: 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 5: 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 5: 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 6: 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 6: 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 |

Table 6: 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) |
|-------------|---|--|--|
| 35 | ge-0/0/35 | xe-0/0/35 | Not supported on this port |
| 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 |

Table 6: 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) |
|-------------|---|--|--|
| 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 |
| 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 7: 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 |

Table 7: 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) |
|-------------|--|--|
| Q3 | xe-0/0/12 | xle-0/1/3 |
| | xe-0/0/13 | |
| | xe-0/0/14 | |
| | xe-0/0/15 | |
| 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 | |

Table 7: 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) |
|-------------|--|--|
| Q9 | xe-0/0/36 xe-0/0/37 xe-0/0/38 xe-0/0/39 | xle-0/1/9 |
| Q10 | xe-0/0/40 xe-0/0/41 xe-0/0/42 xe-0/0/43 | xle-0/1/10 |
| Q11 | xe-0/0/44 xe-0/0/45 xe-0/0/46 xe-0/0/47 | xle-0/1/11 |
| Q12 | xe-0/0/48 xe-0/0/49 xe-0/0/50 xe-0/0/51 | xle-0/1/12 |
| Q13 | xe-0/0/52 xe-0/0/53 xe-0/0/54 xe-0/0/55 | xle-0/1/13 |
| Q14 | xe-0/0/56 xe-0/0/57 xe-0/0/58 xe-0/0/59 | xle-0/1/14 |

Table 7: 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) |
|-------------|--|--|
| Q15 | xe-0/0/60 | xle-0/1/15 |
| | xe-0/0/61 | |
| | xe-0/0/62 | |
| | xe-0/0/63 | |

Table 8: 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 | |

Table 8: 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) |
|-------------|--|--|
| 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 | |
| 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 | |

Table 8: 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) |
|-------------|--|--|
| 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 | |
| Q12 | xe-0/0/12:0 | et-0/0/12 |
| | xe-0/0/12:1 | |
| | xe-0/0/12:2 | |
| | xe-0/0/12:3 | |
| Q13 | xe-0/0/13:0 | et-0/0/13 |
| | xe-0/0/13:1 | |
| | xe-0/0/13:2 | |
| | xe-0/0/13:3 | |
| Q14 | xe-0/0/14:0 | et-0/0/14 |
| | xe-0/0/14:1 | |
| | xe-0/0/14:2 | |
| | xe-0/0/14:3 | |
| Q15 | xe-0/0/15:0 | et-0/0/15 |
| | xe-0/0/15:1 | |
| | xe-0/0/15:2 | |
| | xe-0/0/15:3 | |

Table 9: 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 |
| Q3 | xe-0/0/12 xe-0/0/13 xe-0/0/14 xe-0/0/15 | fte-0/1/3 | xle-0/1/3 |
| Q4 | xe-0/0/16 xe-0/0/17 xe-0/0/18 xe-0/0/19 | fte-0/1/4 | xle-0/1/4 |
| Q5 | xe-0/0/20 xe-0/0/21 xe-0/0/22 xe-0/0/23 | fte-0/1/5 | xle-0/1/5 |
| Q6 | xe-0/0/24 xe-0/0/25 xe-0/0/26 xe-0/0/27 | fte-0/1/6 | xle-0/1/6 |
| Q7 | xe-0/0/28 xe-0/0/29 xe-0/0/30 xe-0/0/31 | fte-0/1/7 | xle-0/1/7 |

Table 9: 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) |
|-------------|--|--|---|
| 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 | | |
| 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 | | |

Table 9: 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) |
|-------------|--|--|---|
| 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 | | |
| Q15 | xe-0/0/60 | Not supported on this port | xle-0/1/15 |
| | xe-0/0/61 | | |
| | xe-0/0/62 | | |
| | xe-0/0/63 | | |

Table 10: 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 |

Table 10: 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) |
|-------------|--|--|
| 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 |
| 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 |

Table 10: 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) |
|-------------|--|--|
| 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 |
| 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 |

Table 10: 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) |
|-------------|--|--|
| 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 |
| 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 11: 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 |

Table 11: 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) |
|-------------|--|---|---|
| 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 |
| 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 |

Table 11: 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) |
|-------------|--|---|---|
| 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 |
| 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 |

Table 11: 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) |
|-------------|--|---|---|
| 48 | Not supported on this port | Not supported on this PIC | fte-0/1/0 <i>NOTE:</i> 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 <i>NOTE:</i> This interface is a fixed fte interface and cannot be changed to xle. |
| 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 12 on page 43](#) 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 12 on page 43](#) for more information regarding the supported system modes for your switch.

Table 12: 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. |

Table 12: System Modes Supported on QFX5100 Switches Running Enhanced Layer 2 Software (*continued*)

| | Default-mode | Mode-104port | Flexi-pic-mode | Non-oversubscribed-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**
- [Interfaces Overview on page 3](#)
 - [Channelizing Interfaces on page 63](#)
 - [Configuring the System Mode on page 77](#)
 - [Understanding Interface Naming Conventions on page 5](#)
 - [Rear Panel of a QFX3500 Device](#)
 - [Front Panel of a QFX3600 Device](#)

Configuring the Port Type on QFX3600 Standalone Switches

The QFX3600 standalone switch provides 16 40-Gbps QSFP+ ports. By default, all 16 ports operate as 40-Gigabit Ethernet (xle) ports. Optionally, you can choose to configure the 40-Gbps ports to operate as four 10-Gigabit Ethernet (xe) ports. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. You can configure up to 64 10-Gigabit Ethernet ports on ports Q0 through Q15.

This topic explains how to configure the port type on QFX3600 standalone switches.



CAUTION: The Packet Forwarding Engine on the QFX3600 standalone switch is restarted when you commit the port type configuration changes. As a result, you might experience packet loss on the switch.

The following message may be displayed in the system log file when the Packet Forwarding Engine is restarted. You can ignore this message.

Pipe write error: Broken pipe

flush operation failed

The following steps describe how to configure either a block of ports or an individual port to operate as 10-Gigabit Ethernet (xe) ports, as well as how to delete a 10-Gigabit Ethernet (xe) port configuration.



NOTE: When you delete the xe port type configuration for an individual port or a block of ports, the ports return to operating as 40-Gigabit Ethernet (xle) ports.

1. To configure a block of ports to operate as 10-Gigabit Ethernet (xe) ports, specify a port range:

```
[edit chassis fpc 0 pic 0]
user@switch# set xe port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q7 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set xe port-range 4 7
```

2. To configure an individual port to operate as a 10-Gigabit Ethernet (xe) port, specify a port number:

```
[edit chassis fpc 0 pic 0]
user@switch# set xe port port-number
```

For example, to configure port Q4 to operate as a 10-Gigabit Ethernet port:

```
[edit chassis fpc 0 pic 0]
user@switch# set xe port 4
```

3. Review your configuration and issue the **commit** command.

```
[edit chassis fpc 0 pic 0]
user@switch# commit
commit complete
```

4. To delete the 10-Gigabit Ethernet (xe) port configuration for a block of ports (and return to the default 40-Gigabit Ethernet configuration), specify a port range:

```
[edit chassis fpc 0 pic 0]
user@switch# delete xe port-range port-range-low port-range-high
```

For example, to delete the 10-Gigabit Ethernet port configuration for ports Q4 through Q7:

```
[edit chassis fpc 0 pic 0]
user@switch# delete xe port-range 4 7
```

5. To delete the 10-Gigabit Ethernet (xe) port configuration for an individual port (and return to the default 40-Gigabit Ethernet configuration), specify a port number:

```
[edit chassis fpc 0 pic 0]
```

```
user@switch# delete xe port port-number
```

For example, to delete the 10-Gigabit Ethernet port configuration for port Q4:

```
[edit chassis fpc 0 pic 0]
user@switch# delete xe port 4
```

Related Documentation

- [Understanding Interface Naming Conventions on page 5](#)
- *pic*

Configuring the QSFP+ Port Type on QFX3500 Standalone Switches

By default, the four 40-Gbps QSFP+ ports are configured to operate as 10-Gigabit Ethernet (xe) ports. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. You can, however, configure the four 40-Gbps QSFP+ ports to operate as 40-Gigabit Ethernet (xle) ports.



NOTE: Port Q0 supports only three (not the typical four) 10-Gigabit Ethernet ports, because one port is reserved.



CAUTION: The Packet Forwarding Engine on the QFX3500 standalone switch is restarted when you commit port type configuration changes (for example, configuring or deleting an xle port). As a result, you might experience packet loss on the device.

The following steps describe how to configure either a block of ports or an individual port to operate as 40-Gigabit Ethernet (xle) ports, as well as how to delete a 40-Gigabit Ethernet (xle) configuration.



NOTE: When you delete an xle block of ports or individual port, the ports return to operating as 10-Gigabit Ethernet ports.

1. To configure a block of ports to operate as 40-Gigabit Ethernet (xle) ports, specify a port range:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port-range port-range-low port-range-high
```

For example, to configure ports Q0 through Q3 to operate as 40-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port-range 0 3
```

2. To configure an individual port to operate as a 40-Gigabit Ethernet (xle) port, specify a port number:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port port-number
```

For example, to configure port Q2 to operate as a 40-Gigabit Ethernet port:

```
[edit chassis fpc 0 pic 2]
user@switch# set xle port 2
```

3. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

4. To delete a block of ports configured as 40-Gigabit Ethernet (xle) ports (and return to the default 10-Gigabit Ethernet configuration), specify a port range:

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port-range port-range-low port-range-high
```

For example, to delete the 40-Gigabit Ethernet (xle) port configuration for ports Q0 through Q3 (and return to the default 10-Gigabit Ethernet configuration):

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port-range 0 3
```

5. To delete an individual port configured as a 40-Gigabit Ethernet (xle) port (and return to the default 10-Gigabit Ethernet configuration), specify an individual port:

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port port-number
```

For example, to delete the 40-Gigabit Ethernet (xle) port configuration for port Q2 (and return to the default 10-Gigabit Ethernet configuration):

```
[edit chassis fpc 0 pic 2]
user@switch# delete xle port 2
```

6. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

- Related Documentation**
- [Understanding Interface Naming Conventions on page 5](#)
 - *pic*

Configuring the QSFP+ Port Type on QFX5100 Devices

You can convert default 40-Gigabit Ethernet data plane uplink interfaces (fte) to 40-Gigabit Ethernet access interfaces (xle) ports, and default 40-Gigabit Ethernet interfaces (xle) to 40-Gigabit Ethernet data plane uplink interfaces (fte). Ports Q0 and Q1 are fixed fte ports and cannot be changed. Ports Q2 and Q3 are fte ports by default but can be changed to xle ports. Ports Q4 and Q5 are xle ports by default but can be changed to fte ports.



NOTE: On QFX5100-24Q switches, ports Q1 through Q7 are fixed FTE ports and cannot be changed.



NOTE: You must configure xle ports in pairs, not individually, otherwise functionality is not guaranteed.



CAUTION: The Packet Forwarding Engine on a QFX5100 switch is restarted when you commit port type configuration changes (for example, configuring or deleting an fte or xle port). As a result, you might experience packet loss on the device.

The following steps describe how to configure either a block of ports or an individual port, as well as how to delete these configurations.

1. To configure a block of ports to operate as 40-Gigabit Ethernet interfaces (xle) , specify a port range:

```
[edit chassis node-group name node-device name pic 1]
user@switch# set xle port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q5 to operate as 40-Gigabit Ethernet interfaces (xle):

```
[edit chassis node-group name node-device name pic 1]
user@switch# set xle port-range 4 5
```

2. To configure a block of ports to operate as 40-Gigabit Ethernet data plane uplink interfaces (fte), specify a port range:

```
[edit chassis node-group name node-device name pic 1]
user@switch# set fte port-range port-range-low port-range-high
```

For example, to configure ports Q4 through Q5 to operate as 40-Gigabit Ethernet data plane uplink interfaces (fte):

```
[edit chassis node-group name node-device name pic 1]
user@switch# set fte port-range 4 5
```

3. To configure an individual port to operate as a 40-Gigabit Ethernet data plane uplink interfaces (fte), specify a port number:

```
[edit chassis node-group name node-device name pic 1]
user@switch# set fte port port-number
```

For example, to configure port Q4 to operate as a 40-Gigabit Ethernet data plane uplink interfaces (fte):

```
[edit chassis node-group name node-device name pic 1]
user@switch# set fte port 4
```

4. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

5. To delete a block of ports configured as 40-Gigabit Ethernet (xle) ports, specify a port range:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port-range port-range-low port-range-high
```

For example, to delete the 40-Gigabit Ethernet access interface (xle) port configuration for ports Q2 through Q3:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port-range 2 3
```

6. To delete an individual port configured as a 40-Gigabit Ethernet (xle) interface:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port port-number
```

For example, to delete the 40-Gigabit Ethernet interface (xle) for port Q2:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete xle port 2
```

7. To delete a block of ports configured as 40-Gigabit Ethernet data plane uplink interfaces (fte), specify a port range:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port-range port-range-low port-range-high
```

For example, to delete the block of ports configured as 40-Gigabit Ethernet data plane uplink interfaces (fte) for ports Q4 through Q5:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port-range 4 5
```

8. To delete an individual port configured as a 40-Gigabit Ethernet data plane uplink interfaces (fte):

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port port-number
```

For example, to delete the 40-Gigabit Ethernet data plane uplink interfaces (fte) for port Q4:

```
[edit chassis node-group name node-device name pic 1]
user@switch# delete fte port 4
```

9. Review your configuration and issue the **commit** command.

```
[edit]  
user@switch# commit  
commit complete
```

**Related
Documentation**

- [Understanding Interface Naming Conventions on page 5](#)
- [Understanding Port Ranges and System Modes on page 18](#)
- *pic*

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

```
[edit interfaces ]
```

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



NOTE:

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

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

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



NOTE: The broadcast address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255

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

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

4. [Optional] For interfaces that carry IP version 6 (IPv6) traffic, configure the host to assign 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

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



NOTE: RSTP and LLDP are not supported on the OCX Series.

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 54](#)
- [Configuring the Link Settings for Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-96S, and EX4600 Switches on page 55](#)
- [Configuring Gigabit Ethernet Interfaces on QFX5100-48T Switches on page 55](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48S, QFX5100-24Q, QFX5100-96S, and EX4600 Switches on page 57](#)
- [Configuring the Link Settings for 10-Gigabit Ethernet Interfaces on QFX5100-48T Switches on page 57](#)
- [Configuring the IP Options on QFX5100-48S, QFX5100-48T, QFX5100-24Q, and EX4600 Switches on page 58](#)

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 original CLI:

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

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
```

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.

- 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, 1G, and 10G.

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.

- 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.

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 reenable autonegotiation 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 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.

- 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.

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.
- The speed cannot be configured.
- 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
```

Related Documentation

- [Monitoring Interface Status and Traffic on page 79](#)
- [show interfaces xe on page 409](#)
- [show interfaces ge-](#)
- [speed on page 266](#)
- [Understanding Interface Naming Conventions on page 5](#)

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 on page 53](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)

[Configuring an LPM Table With Junos OS Release 13.2X51-D10](#)

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

[edit]

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

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



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

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

[Table 13 on page 61](#) provides examples of valid combinations that the LPM table can store, also using the **l2-profile-one** profile. Once again, each row in the table represents a case in which the table is full and cannot accommodate any more entries.

Table 13: Example LPM Table Combinations Using l2-profile-one With Junos OS 13.2X51-D10

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

[Table 14 on page 62](#) provides examples of valid combinations that the LPM table can store when you use the **lpm-profile** profile. As before, each row represents a case in which the table is full and cannot accommodate any more entries.

Table 14: Example LPM Table Combinations Using lpm-profile With Junos OS 13.2X51-D10

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

Related Documentation

- *Configuring the Unified Forwarding Table*

Channelizing Interfaces



NOTE: On QFX10008 and QFX10016 switches, channelization is supported on fiber break-out cables using standard structured cabling techniques. Channelization is not supported on the QFX10000-30C line card.



NOTE: On QFX10002 switches running on Junos OS Release 15.1X53-D10 or Junos OS Release 15.1X53-D15, when you delete and then reapply channelized interfaces, traffic is disrupted and might not be recovered.

The QFX3500, QFX3600, QFX5100, and EX4600, QFX10002, QFX10008, and QFX10016 switches provide 40-Gbps QSFP+ ports that can be channelized. Channelization allows you to configure 40-Gbps QSFP+ ports to operate as four 10-Gigabit Ethernet (xe) interfaces. You can use QSFP+ to four SFP+ breakout cables or QSFP+ transceivers with fiber breakout cables to connect the 10-Gigabit Ethernet ports to other servers, storage, and switches. By default, the four 40-Gbps QSFP+ ports operate as 40-Gigabit Ethernet (et) ports. When an et port is channelized to four xe ports, a colon is used to signify the four separate channels. For example, on a switch with port 2 on PIC 1 configured as four 10-Gigabit Ethernet ports, the interface names are *xe-0/1/2:0*, *xe-0/1/2:1*, *xe-0/1/2:2*, and *xe-0/1/2:3*.

By default, the 40-Gbps QSFP+ ports on EX4600 and QFX5100 switches are channelized automatically (auto-channelized) if any of the four channels on a 40-Gbps QSFP+ port receive data, unless you have configured channelization either at the chassis level or at the port level. Auto-channelization is not supported on interfaces contained in expansion modules, or on Virtual Chassis ports.

You can disable auto-channelization by including the **disable-auto-speed-detection** statement at the **[edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high) channel-speed]** hierarchy.

There are restrictions on the ports you can channelize on the QFX5100-24Q and QFX5100-96S switches, depending on the system mode you enable. If you try to channelize ports that are restricted, the configuration is ignored. See [“Configuring the System Mode” on page 77](#) for more information.

On QFX10002, QFX10008, and QFX10016 switches, there are 100-Gigabit Ethernet ports that work either as 100-Gigabit Ethernet or as 40-Gigabit Ethernet, but are recognized as 40-Gigabit Ethernet by default. You cannot channelize the 100-Gigabit Ethernet ports when they are operating as 100-Gigabit Ethernet interfaces. The 40-Gigabit Ethernet ports can operate independently or be channelized into four 10-Gigabit Ethernet ports as part of a port range. Ports cannot be channelized individually. Only the first and fourth port in each 6XQSFP cage is available to channelize as part of a port range. In a port range, the ports are bundled with the next two consecutive ports. For example, if you want to channelize ports 0 through 2, you would channelize port 0 only. If you try to channelize a port that is not supported, you will receive an error message when you commit the configuration. Auto-channelization is not supported on any ports.

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.

[Table 15 on page 65](#) provides detailed information on which ports are 100-Gigabit Ethernet, which ports can be channelized, and which ports are disabled when a 100-Gigabit Ethernet is inserted in the QFX10002-36Q switch and the QFX10000-36Q line card on a QFX10008 or QFX10016 switch. [Table 16 on page 67](#) provides detailed information on which ports are 100-Gigabit Ethernet, which ports can be channelized, and which ports are disabled when a 100-Gigabit Ethernet is inserted in the QFX10002-72Q switch. On the QFX10008 and QFX10016 switches with the QFX10000-36Q line card installed, only ports 0 through 35 are available. For more information, see *QFX10002-72Q Port Panel* and *QFX10000-36Q Line Card*.

Table 15: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings

| Port Number | 4X10 Gigabit Ethernet Port | 4X10 Gigabit Channelized Port Group | 40-Gigabit Ethernet (Default) | 100-Gigabit Ethernet | 100-Gigabit Ethernet Disables |
|-------------|----------------------------|-------------------------------------|-------------------------------|----------------------|-------------------------------|
| 0 | ✓ | ✓ | ✓ | – | – |
| 1 | ✓ | | ✓ | ✓ | 0, 2 |
| 2 | ✓ | | ✓ | – | – |
| 3 | ✓ | ✓ | ✓ | – | – |
| 4 | ✓ | | ✓ | – | – |
| 5 | ✓ | | ✓ | ✓ | 3, 4 |
| 6 | ✓ | ✓ | ✓ | – | – |
| 7 | ✓ | | ✓ | ✓ | 6, 8 |
| 8 | ✓ | | ✓ | – | – |
| 9 | ✓ | ✓ | ✓ | – | – |
| 10 | ✓ | | ✓ | – | – |
| 11 | ✓ | | ✓ | ✓ | 9, 10 |
| 12 | ✓ | ✓ | ✓ | – | – |
| 13 | ✓ | | ✓ | ✓ | 12, 14 |
| 14 | ✓ | | ✓ | – | – |

Table 15: QFX10002-36Q Switch and QFX10000-36Q Line Card Port Mappings (*continued*)

| Port Number | 4X10 Gigabit Ethernet Port | 4X10 Gigabit Channelized Port Group | 40-Gigabit Ethernet (Default) | 100-Gigabit Ethernet | 100-Gigabit Ethernet Disables |
|-------------|----------------------------|-------------------------------------|-------------------------------|----------------------|-------------------------------|
| 15 | ✓ | ✓ | ✓ | – | – |
| 16 | ✓ | | ✓ | – | – |
| 17 | ✓ | | ✓ | ✓ | 15, 16 |
| 18 | ✓ | ✓ | ✓ | – | – |
| 19 | ✓ | | ✓ | ✓ | 18, 20 |
| 20 | ✓ | | ✓ | – | – |
| 21 | ✓ | ✓ | ✓ | – | – |
| 22 | ✓ | | ✓ | – | – |
| 23 | ✓ | | ✓ | ✓ | 21, 22 |
| 24 | ✓ | ✓ | ✓ | – | – |
| 25 | ✓ | | ✓ | ✓ | 24, 26 |
| 26 | ✓ | | ✓ | – | – |
| 27 | ✓ | ✓ | ✓ | – | – |
| 28 | ✓ | | ✓ | – | – |
| 29 | ✓ | | ✓ | ✓ | 27, 28 |
| 30 | ✓ | ✓ | ✓ | – | – |
| 31 | ✓ | | ✓ | ✓ | 30, 32 |
| 32 | ✓ | | ✓ | – | – |
| 33 | ✓ | ✓ | ✓ | – | – |
| 34 | ✓ | | ✓ | – | – |
| 35 | ✓ | | ✓ | ✓ | 33, 34 |

Table 16: QFX10002-72Q Switch Port Mappings

| Port Number | 4X10 Gigabit Ethernet Port | 4X10 Gigabit Channelized Port Group | 40-Gigabit Ethernet (Default) | 100-Gigabit Ethernet | 100-Gigabit Ethernet Disables |
|-------------|----------------------------|-------------------------------------|-------------------------------|----------------------|-------------------------------|
| 0 | ✓ | ✓ | ✓ | – | – |
| 1 | ✓ | | ✓ | ✓ | 0, 2 |
| 2 | ✓ | | ✓ | – | – |
| 3 | ✓ | ✓ | ✓ | – | – |
| 4 | ✓ | | ✓ | – | – |
| 5 | ✓ | | ✓ | ✓ | 3, 4 |
| 6 | ✓ | ✓ | ✓ | – | – |
| 7 | ✓ | | ✓ | ✓ | 6, 8 |
| 8 | ✓ | | ✓ | – | – |
| 9 | ✓ | ✓ | ✓ | – | – |
| 10 | ✓ | | ✓ | – | – |
| 11 | ✓ | | ✓ | ✓ | 9, 10 |
| 12 | ✓ | ✓ | ✓ | – | – |
| 13 | ✓ | | ✓ | ✓ | 12, 14 |
| 14 | ✓ | | ✓ | – | – |
| 15 | ✓ | ✓ | ✓ | – | – |
| 16 | ✓ | | ✓ | – | – |
| 17 | ✓ | | ✓ | ✓ | 15, 16 |
| 18 | ✓ | ✓ | ✓ | – | – |
| 19 | ✓ | | ✓ | ✓ | 18, 20 |
| 20 | ✓ | | ✓ | – | – |

Table 16: QFX10002-72Q Switch Port Mappings (*continued*)

| Port Number | 4X10 Gigabit Ethernet Port | 4X10 Gigabit Channelized Port Group | 40-Gigabit Ethernet (Default) | 100-Gigabit Ethernet | 100-Gigabit Ethernet Disables |
|-------------|----------------------------|-------------------------------------|-------------------------------|----------------------|-------------------------------|
| 21 | ✓ | ✓ | ✓ | – | – |
| 22 | ✓ | | ✓ | – | – |
| 23 | ✓ | | ✓ | ✓ | 21, 22 |
| 24 | ✓ | ✓ | ✓ | – | – |
| 25 | ✓ | | ✓ | ✓ | 24, 26 |
| 26 | ✓ | | ✓ | – | – |
| 27 | ✓ | ✓ | ✓ | – | – |
| 28 | ✓ | | ✓ | – | – |
| 29 | ✓ | | ✓ | ✓ | 27, 28 |
| 30 | ✓ | ✓ | ✓ | – | – |
| 31 | ✓ | | ✓ | ✓ | 30, 32 |
| 32 | ✓ | | ✓ | – | – |
| 33 | ✓ | ✓ | ✓ | – | – |
| 34 | ✓ | | ✓ | – | – |
| 35 | ✓ | | ✓ | ✓ | 33, 34 |
| 36 | ✓ | ✓ | ✓ | – | – |
| 37 | ✓ | | ✓ | ✓ | 36, 38 |
| 38 | ✓ | | ✓ | – | – |
| 39 | ✓ | ✓ | ✓ | – | – |
| 40 | ✓ | | ✓ | – | – |
| 41 | ✓ | | ✓ | ✓ | 39, 40 |

Table 16: QFX10002-72Q Switch Port Mappings (*continued*)

| Port Number | 4X10 Gigabit Ethernet Port | 4X10 Gigabit Channelized Port Group | 40-Gigabit Ethernet (Default) | 100-Gigabit Ethernet | 100-Gigabit Ethernet Disables |
|-------------|----------------------------|-------------------------------------|-------------------------------|----------------------|-------------------------------|
| 42 | ✓ | ✓ | ✓ | – | – |
| 43 | ✓ | | ✓ | ✓ | 42, 44 |
| 44 | ✓ | | ✓ | – | – |
| 45 | ✓ | ✓ | ✓ | – | – |
| 46 | ✓ | | ✓ | – | – |
| 47 | ✓ | | ✓ | ✓ | 45, 46 |
| 48 | ✓ | ✓ | ✓ | – | – |
| 49 | ✓ | | ✓ | ✓ | 48, 50 |
| 50 | ✓ | | ✓ | – | – |
| 51 | ✓ | ✓ | ✓ | – | – |
| 52 | ✓ | | ✓ | – | – |
| 53 | ✓ | | ✓ | ✓ | 51, 52 |
| 54 | ✓ | ✓ | ✓ | – | – |
| 55 | ✓ | | ✓ | ✓ | 54, 56 |
| 56 | ✓ | | ✓ | – | – |
| 57 | ✓ | ✓ | ✓ | – | – |
| 58 | ✓ | | ✓ | – | – |
| 59 | ✓ | | ✓ | ✓ | 57, 58 |
| 60 | ✓ | ✓ | ✓ | – | – |
| 61 | ✓ | | ✓ | ✓ | 60, 62 |
| 62 | ✓ | | ✓ | – | – |

Table 16: QFX10002-72Q Switch Port Mappings (*continued*)

| Port Number | 4X10 Gigabit Ethernet Port | 4X10 Gigabit Channelized Port Group | 40-Gigabit Ethernet (Default) | 100-Gigabit Ethernet | 100-Gigabit Ethernet Disables |
|-------------|----------------------------|-------------------------------------|-------------------------------|----------------------|-------------------------------|
| 63 | ✓ | ✓ | ✓ | – | – |
| 64 | ✓ | | ✓ | – | – |
| 65 | ✓ | | ✓ | ✓ | 63, 64 |
| 66 | ✓ | ✓ | ✓ | – | – |
| 67 | ✓ | | ✓ | ✓ | 66, 68 |
| 68 | ✓ | | ✓ | – | – |
| 69 | ✓ | ✓ | ✓ | – | – |
| 70 | ✓ | | ✓ | – | – |
| 71 | ✓ | | ✓ | ✓ | 69, 70 |

The following steps describe how to configure a block of ports or an individual port to operate as 10-Gigabit Ethernet ports.

1. To configure a block of 40-Gigabit Ethernet (*et*) ports on QFX3500, QFX3600, QFX5100, EX4600 switches to operate as 10-Gigabit Ethernet ports, specify a port range and channel speed:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port-range port-range-low port-range-high channel-speed speed
```

For example, to configure ports 0 through 3 on PIC 1 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 1]
user@switch# set port-range 0 3 channel-speed 10g
```

2. To configure a range of ports on a QFX10002-72Q, QFX10008, or QFX10016 switch to operate as 10-Gigabit Ethernet ports:



NOTE: The `port-range` statement is not available on QFX10002-72Q, QFX10008, and QFX10016 switches. Instead, configure the port range using the `port` statement. Starting from port 0, you channelize every third port to channelize a group of three ports. For example, channelize port 0 to channelize ports 0 through 2, port 3 to channelize ports 3 through 5, and so on. See [Table 15 on page 65](#) for port mapping information.

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port port-number channel-speed speed
```

For example, to configure ports 0 through 2 on PIC 0 to operate as 10-Gigabit Ethernet ports:



NOTE: When you channelize port 0, ports 1 and 2 are also channelized.

```
[edit chassis fpc 0 pic 1]
user@switch# set port 0 channel-speed 10g
```

3. To configure an individual 40-Gigabit Ethernet (*et*) port on QFX3500, QFX3600, QFX5100, and EX4600 switches to operate as 10-Gigabit Ethernet (*xe*) ports, specify a port number and channel speed:

```
[edit chassis fpc 0 pic 0]
user@switch# set port port-number channel-speed speed
```

For example, to configure port 3 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port 3 channel-speed 10g
```

4. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

5. To return a range of ports on QFX3500, QFX3600, QFX5100, and EX4600 switches to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-range port-range-low port-range-high channel-speed speed
```

For example, to return ports 0 through 3 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-range 0 3 channel-speed 10g
```

6. To return a range of ports on QFX10002-72Q, QFX10008, or QFX10016 switches to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port port-number channel-speed speed
```

For example, to return ports 0 through 2 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 1]
user@switch# delete port-0 channel-speed 10g
```

7. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

8. To return a port on QFX3500, QFX3600, QFX5100, and EX4600 switches to the default 40-Gigabit Ethernet configuration, delete the 10g statement:

```
[edit chassis fpc 0 pic 0]
```

```
user@switch# delete port port-number channel-speed speed
```

For example, to return port 2 to the default 40-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]  
user@switch# delete port 2 channel-speed 10g
```

9. Review your configuration and issue the **commit** command.

```
[edit]  
user@switch# commit  
commit complete
```

The following steps describe how to disable auto-channelization at the port level on QFX3500, QFX3600, QFX5100, and EX4600 switches.

1. To disable auto-channelization at the port level, include the **disable** statement:

```
[edit]  
user@switch# set chassis fpc slot-number pic pic-number (port port-number |  
    port-range port-range-low port-range-high) channel-speed  
    disable-auto-speed-detection
```

For example, to disable auto-channelization for one port:

```
[edit]  
user@switch# set chassis fpc 0 pic 0 port 2 channel-speed  
    disable-auto-speed-detection
```

For example, to disable auto-channelization for a range of ports:

```
[edit]  
user@switch# set chassis fpc 0 pic 0 port-range 2 4 channel-speed  
    disable-auto-speed-detection
```

2. Review your configuration and issue the **commit** command.

```
[edit]  
user@switch# commit  
commit complete
```

Related Documentation

- [Configuring the System Mode on page 77](#)
- [channel-speed on page 218](#)
- [fpc on page 234](#)
- [pic on page 263](#)

Channelizing Interfaces on QFX5110-48S Switches

On the QFX5110-48S switch, there are four ports labeled 48 through 51, which support QSFP28 ports. The QSFP28 ports support 100-Gigabit Ethernet interfaces and 40-Gigabit Ethernet interfaces. You can channelize the 40-Gigabit Ethernet interfaces to four independent 10-Gigabit Ethernet interfaces using breakout cables.

When you channelize the 40-Gigabit Ethernet interfaces as 10-Gigabit Ethernet interfaces, the interface names appear in the `xe-fpc/pic/port:channel` format, where `channel` can be a value of 0 through 3. To channelize the ports, manually configure the port speed using the **`set chassis fpc slot-number port port-number channel-speed speed`** command, where the speed can be set to 10G. The ports do not support auto-channelization.



NOTE: On QFX5110-48S standalone switches, the FPC value is always 0. Also, Virtual Chassis is not supported.

The following steps describe how to channelize blocks of ports or individual ports.

1. To configure a block of 100-Gigabit Ethernet (*et*) ports to operate as 40-Gigabit Ethernet ports, specify a port range and channel speed:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port-range port-range-low port-range-high channel-speed speed
```

For example, to configure ports 48 through 51 on PIC 0 to operate as 40-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port-range 48-51 channel-speed 40g
```

2. To configure an individual 40-Gigabit Ethernet (*et*) port to operate as 10-Gigabit Ethernet (*xe*) ports, specify a port number and channel speed:

```
[edit chassis fpc 0 pic 0]
user@switch# set port port-number channel-speed speed
```

For example, to configure port 48 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port 48 channel-speed 10g
```

3. Review your configuration and issue the **`commit`** command.

```
[edit]
user@switch# commit
commit complete
```

4. To return a range of ports from the 40-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration, delete the 40g statement:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port-range port-range-low port-range-high channel-speed speed
```

For example, to return ports 48 through 51 from the 40-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port-range 48-51 channel-speed 40g
```

5. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

6. To return an individual 40-Gigabit Ethernet port to the default 100-Gigabit Ethernet configuration, delete the 40g statement:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port port-number channel-speed speed
```

For example, to return port 48 from the 40-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port 48 channel-speed 40g
```

7. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

- Related Documentation**
- [channel-speed on page 218](#)
 - [fpc on page 234](#)
 - [pic on page 263](#)

Channelizing Interfaces on QFX5200 Switches

You can channelize the 100-Gigabit Ethernet interfaces to two independent 50-Gigabit Ethernet or to four independent 25-Gigabit Ethernet interfaces. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces, and in this configuration can either operate as dedicated 40-Gigabit Ethernet interfaces or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables.

There are a total of 32 physical ports on the QFX5200 switch. Any port can be used as either 100-Gigabit Ethernet or 40-Gigabit Ethernet interfaces. You choose the speed by plugging in the appropriate transceiver. They can also be channelized to 50G, 25G or 10G.

By default, the 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces appear in the `et-fpc/pic/port` format. When the 100-Gigabit Ethernet interfaces are channelized as 50-Gigabit Ethernet and 25-Gigabit Ethernet interfaces, the interface names appear in the `et-fpc/pic/port:channel` format. When the 40-Gigabit Ethernet interfaces are channelized as 10-Gigabit Ethernet interfaces, the interface names appear in the `xe-fpc/pic/port:channel` format, where channel can be a value of 0 through 3. To channelize the ports, manually configure the port speed using the **`set chassis fpc slot-number port port-number channel-speed speed`** command, where the speed can be set to 10G, 25G, or 50G. If a 100-Gigabit Ethernet transceiver is connected, you can only set the speed to 25G or 50G. If a 40-Gigabit Ethernet transceiver is connected, you can only set the speed to 10G. There is no commit check for this, however. The ports do not support autochannelization.



NOTE: For details about supported transceivers and cable specifications, see the [QFX5200 Switch Hardware Guide](#).



NOTE: On QFX5200 standalone switches, the FPC value is always 0. Virtual Chassis is not supported.

The following steps describe how to channelize blocks of ports or individual ports.

1. To configure a block of 100-Gigabit Ethernet (*et*) ports to operate as 50-Gigabit Ethernet ports, specify a port range and channel speed:

```
[edit chassis fpc fpc-slot pic pic-slot]
user@switch# set port-range port-range-low port-range-high channel-speed speed
```

For example, to configure ports 0 through 3 on PIC 0 to operate as 50-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port-range 0 3 channel-speed 50g
```

2. To configure a block of 100-Gigabit Ethernet (*et*) ports to operate as 25-Gigabit Ethernet ports, specify a port range and channel speed:

```
[edit chassis fpc fpc-slot pic pic-slot]
```

```
user@switch# set port-range port-range-low port-range-high channel-speed speed
```

For example, to configure ports 0 through 3 on PIC 0 to operate as 25-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port-range 0 3 channel-speed 25g
```

3. To configure an individual 40-Gigabit Ethernet (*et*) port to operate as 10-Gigabit Ethernet (*xe*) ports, specify a port number and channel speed:

```
[edit chassis fpc 0 pic 0]
user@switch# set port port-number channel-speed speed
```

For example, to configure port 3 to operate as 10-Gigabit Ethernet ports:

```
[edit chassis fpc 0 pic 0]
user@switch# set port 3 channel-speed 10g
```

4. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

5. To return a range of ports from the 50-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration, delete the 50g statement:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port-range port-range-low port-range-high channel-speed speed
```

For example, to return ports 0 through 3 from the 50-Gigabit Ethernet configuration to the default 100-Gigabit Ethernet configuration:

```
[edit chassis fpc 0 pic 0]
user@switch# delete port-range 0 3 channel-speed 50g
```



NOTE: To configure the ports to another channel-speed, you must delete the current port-range statement to return to the default 100-Gigabit Ethernet configuration.

6. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

7. Review your configuration and issue the **commit** command.

```
[edit]
user@switch# commit
commit complete
```

Related Documentation

- [channel-speed on page 218](#)
- [fpc on page 234](#)
- [pic on page 263](#)

Configuring the System Mode

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. By default, all QSFP+ interfaces are auto-channelized. Auto-channelization is not supported on interfaces contained in expansion modules or on Virtual Chassis ports. To disable auto-channelization, see “Channelizing Interfaces” on page 63 for more information.



NOTE: The QFX5200 switches do not support System Mode.



NOTE: When you request the system mode change, we recommend that you reboot the switch for the system mode to take effect.



CAUTION: The Packet Forwarding Engine on the switch is restarted when you issue system mode changes. As a result, you might experience packet loss on the switch.

See [Table 17 on page 77](#), [Table 18 on page 78](#), and [Table 19 on page 79](#) for more information regarding the supported system modes for your switch.

Table 17: System Modes Supported on QFX5100 Switches with QFX-EM-4Q or QFX-PFA-4Q Expansion Modules Installed

| | Default-mode | Mode-104port | Flexi-pic-mode | Non-oversubscribed-mode |
|-------------|---------------|---------------|----------------|-------------------------|
| QFX5100-48S | Not supported | Not supported | Not supported | Not supported |

Table 17: System Modes Supported on QFX5100 Switches with QFX-EM-4Q or QFX-PFA-4Q Expansion Modules Installed (*continued*)

| | Default-mode | Mode-104port | Flexi-pic-mode | Non-oversubscribed-mode |
|-------------|---|---|--|--|
| QFX5100-24Q | <p>Supported</p> <p>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.</p> | <p>Supported</p> <p>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.</p> | <p>Supported</p> <p>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.</p> | <p>Supported</p> <p>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.</p> |
| QFX5100-96S | <p>Supported</p> <p>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.</p> | <p>Not supported</p> | <p>Not supported</p> | <p>Supported</p> <p>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.</p> |

Table 18: System Modes Supported on QFX5100-24Q Switches with the EX4600-8F Expansion Module Installed

| | Default-mode | Mode-104port | Flexi-pic-mode | Non-oversubscribed-mode |
|-------------|----------------------|----------------------|--|--|
| QFX5100-24Q | <p>Not supported</p> | <p>Not supported</p> | <p>Supported</p> <p>On PIC 0, you cannot channelize ports 0 through 3.</p> | <p>Not supported</p> <p>Expansion modules cannot be installed in PICs 1 and 2.</p> |

Table 19: System Modes Supported on QFX5100-24Q Switches with EX4600-8F and QFX-EM-4Q Expansion Modules Installed

| | Default-mode | Mode-104port | Flexi-pic-mode | Non-oversubscribed-mode |
|-------------|---|--|--|--|
| QFX5100-24Q | Only the QFX-EM-4Q module is supported. | <p>Only the QFX-EM-4Q module is supported.</p> <p>If you have installed the EX4600-8F expansion module on PIC 1, and you have installed the QFX-EM-4Q module on PIC 2, The 40-Gbps QSFP+ ports in both PIC slots are not supported.</p> <p>If you have installed the EX4600-8F expansion module on PIC 2, and you have installed the QFX-EM-4Q module on PIC 1, only the QFX-EM-4Q module on PIC 1 is supported.</p> | <p>Supported</p> <p>On PIC 0, you cannot channelize ports 0 through 3.</p> | <p>Not supported</p> <p>You cannot install the QFX-EM-4Q or EX4600-8F modules on PICs 1 and 2.</p> |

The following steps describe how to change the system mode.

1. To change the system mode, issue the following operational command:

```
{master:0}
root> request chassis system-mode mode
```

For example:

```
{master:0}
root> request chassis system-mode non-oversubscribed-mode
```

2. To return to the default mode (default-mode), issue the following operational command:

```
{master:0}
root> request chassis system-mode default-mode
```

3. To see which system mode is configured, issue the following operational command:

```
{master:0}
root> show chassis system-mode
```

Related Documentation

- [Understanding Interface Naming Conventions on page 5](#)
- [Understanding Port Ranges and System Modes on page 18](#)
- [Channelizing Interfaces on page 63](#)

Monitoring Interface Status and Traffic

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

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

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

Troubleshooting Network Interfaces

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

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

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

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

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

PART 2

Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management on page 83](#)
- [Ethernet OAM Link Fault Management Operational Commands on page 91](#)

CHAPTER 2

Understanding Ethernet OAM Link Fault Management

- [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 83](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 84](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86](#)

Understanding Ethernet OAM Link Fault Management for an EX Series Switch

Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches 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 on EX Series switches:

- 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 EX Series 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.

**Related
Documentation**

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 84](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86](#)

Configuring Ethernet OAM Link Fault Management (CLI Procedure)

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
```

Related Documentation

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86](#)
- [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 83](#)

Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches

Junos OS for EX Series switches 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.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet interface:

- [Requirements on page 87](#)
- [Overview and Topology on page 87](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 1 on page 87](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 2 on page 88](#)
- [Verification on page 89](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.4 or later for EX Series switches
- Two EX3200 or EX4200 switches connected directly

Overview and Topology

Junos OS for EX Series switches 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.

This example uses two EX4200 switches connected directly. Before you begin configuring Ethernet OAM LFM on two switches, connect the two switches directly through a trunk interface.

Configuring Ethernet OAM Link Fault Management on Switch 1

CLI Quick Configuration To quickly configure Ethernet OAM LFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/0
set interface ge-0/0/0 link-discovery active
set interface ge-0/0/0 pdu-interval 800
set interface ge-0/0/0 remote-loopback
```

Step-by-Step Procedure To configure Ethernet OAM LFM on switch 1:

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


```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface (OAM LFM) ge-0/0/0
```
2. Specify that the interface initiates the discovery process by configuring the link discovery mode to **active**:


```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface ge-0/0/0 link-discovery active
```
3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800 on switch 1:


```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface pdu-interval 800
```

4. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Ensure that the remote DTE supports remote loopback mode. To set the remote DTE in loopback mode

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface ge-0/0/0.0 remote-loopback
```

Results

Check the results of the configuration:

```
[edit]
user@switch1# show

protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/0 {
          pdu-interval 800;
          link-discovery active;
          remote-loopback;
        }
      }
    }
  }
}
```

Configuring Ethernet OAM Link Fault Management on Switch 2

CLI Quick Configuration To quickly configure Ethernet OAM LFM on switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management ]
set interface ge-0/0/1
set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

Step-by-Step Procedure To configure Ethernet OAM LFM on switch 2:

1. Enable OAM on the peer interface on switch 2:

```
[edit protocols oam ethernet link-fault-management]
user@switch2# set interface ge-0/0/1
```
2. Enable remote loopback support for the local interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

Results Check the results of the configuration:

```
[edit]
user@switch2# show

protocols {
  oam {
    ethernet {
      link-fault-management {
        interface ge-0/0/1 {
          negotiation-options {
```

```

    allow-remote-loopback;
  }
}
}
}
}

```

Verification

Verifying That OAM LFM Has Been Configured Properly

- Purpose** Verify that OAM LFM has been configured properly.
- Action** Use the `show oam ethernet link-fault-management` command:
- ```
user@switch1#show oam ethernet link-fault-management
```

## Sample Output

```

Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
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

```

- Meaning** When the output displays the MAC address and the discover state is **Send Any**, it means that OAM LFM has been configured properly.

- Related Documentation**
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 84](#)
  - [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 83](#)





## CHAPTER 3

# Ethernet OAM Link Fault Management Operational Commands

- `show oam ethernet link-fault-management`

## show oam ethernet link-fault-management

|                                 |                                                                                                                                                                                                                                                                   |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | show oam ethernet link-fault-management<br><brief   detail><br><interface-name>                                                                                                                                                                                   |
| <b>Release Information</b>      | Command introduced in Junos OS Release 9.4 for EX Series switches.                                                                                                                                                                                                |
| <b>Description</b>              | Displays Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.                                                                                                                                        |
| <b>Options</b>                  | <b>brief   detail</b> —(Optional) Display the specified level of output.<br><br><b>interface-name</b> —(Optional) Display link fault management information for the specified Ethernet interface only.                                                            |
| <b>Required Privilege Level</b> | view                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86</a></li> <li>• <a href="#">Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84</a></li> </ul> |
| <b>List of Sample Output</b>    | <a href="#">show oam ethernet link-fault-management brief on page 96</a><br><a href="#">show oam ethernet link-fault-management detail on page 96</a>                                                                                                             |
| <b>Output Fields</b>            | Table 20 on page 92 lists the output fields for the <b>show oam ethernet link-fault-management</b> command. Output fields are listed in the approximate order in which they appear.                                                                               |

Table 20: show oam ethernet link-fault-management Output Fields

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

Table 20: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                       | Field Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Level of Output |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Flags</b>                     | Information about the interface. <ul style="list-style-type: none"> <li>• <b>Remote-Stable</b>—Indicates remote OAM client acknowledgment of, and satisfaction with local OAM state information. <b>False</b> indicates that remote DTE has either not seen or is unsatisfied with local state information. <b>True</b> indicates that remote DTE has seen and is satisfied with local state information.</li> <li>• <b>Local-Stable</b>—Indicates local OAM client acknowledgment of, and satisfaction with remote OAM state information. <b>False</b> indicates that local DTE either has not seen or is unsatisfied with remote state information. <b>True</b> indicates that local DTE has seen and is satisfied with remote state information.</li> <li>• <b>Remote-State-Valid</b>—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. <b>False</b> indicates that OAM client has not seen remote state information. <b>True</b> indicates that the OAM client has seen remote state information.</li> </ul>                                                                              | All levels      |
| <b>Remote loopback status</b>    | 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      |
| <b>Remote entity information</b> | Remote entity information. <ul style="list-style-type: none"> <li>• <b>Remote MUX action</b>—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.</li> <li>• <b>Remote parser action</b>—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.</li> <li>• <b>Discovery mode</b>—Indicates whether discovery mode is active or inactive.</li> <li>• <b>Unidirectional mode</b>—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes.</li> <li>• <b>Remote loopback mode</b>—Indicates whether remote loopback is supported or not supported.</li> <li>• <b>Link events</b>—Indicates whether interpreting link events is supported or not supported on the remote peer.</li> <li>• <b>Variable requests</b>—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.</li> </ul> | All levels      |
| <b>OAM Receive Statistics</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                 |
| <b>Information</b>               | The number of information PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>detail</b>   |
| <b>Event</b>                     | The number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>   |
| <b>Variable request</b>          | The number of variable request PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>   |
| <b>Variable response</b>         | The number of variable response PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>detail</b>   |
| <b>Loopback control</b>          | The number of loopback control PDUs received.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>detail</b>   |

Table 20: show oam ethernet link-fault-management Output Fields (*continued*)

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

Table 20: show oam ethernet link-fault-management Output Fields (*continued*)

| Field Name                                               | Field Description                                                                                                                                                                       | Level of Output |
|----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Total errors</b>                                      | The number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset.<br><br>A frame error is any frame error on the underlying physical layer. | <b>detail</b>   |
| <b>OAM Received Frame Period Error Event Information</b> |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.                                                                                 | <b>detail</b>   |
| <b>Window</b>                                            | The duration of the frame seconds window.                                                                                                                                               | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of frame seconds errors in the period.                                                                                                                                       | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.                                                                     | <b>detail</b>   |
| <b>OAM Transmitted Symbol Error Event Information</b>    |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.                                                                                      | <b>detail</b>   |
| <b>Window</b>                                            | The symbol error event window in the transmitted PDU.                                                                                                                                   | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of errored symbols in the period required for the event to be generated.                                                                                                     | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of symbol errors in the period reported in the transmitted event PDU.                                                                                                        | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.                                                                       | <b>detail</b>   |
| <b>OAM Transmitted Frame Error Event Information</b>     |                                                                                                                                                                                         |                 |
| <b>Events</b>                                            | The number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.                                                                                     | <b>detail</b>   |
| <b>Window</b>                                            | The duration of the window in terms of the number of 100 ms period intervals.                                                                                                           | <b>detail</b>   |
| <b>Threshold</b>                                         | The number of detected errored frames required for the event to be generated.                                                                                                           | <b>detail</b>   |
| <b>Errors in period</b>                                  | The number of detected errored frames in the period.                                                                                                                                    | <b>detail</b>   |
| <b>Total errors</b>                                      | The number of errored frames that have been detected after the OAM sublayer was reset.                                                                                                  | <b>detail</b>   |

## 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
```

## PART 3

# Generic Routing Encapsulation (GRE)

- [Understanding GRE on page 99](#)





## CHAPTER 4

# Understanding GRE

- [Understanding Generic Routing Encapsulation on page 99](#)
- [Configuring Generic Routing Encapsulation Tunneling on page 102](#)
- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 103](#)

## Understanding Generic Routing Encapsulation

---

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

This topic describes:

- [Overview of GRE on page 99](#)
- [GRE Tunneling on page 100](#)
- [Using a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 and OCX Series Switches on page 101](#)
- [Configuration Limitations on page 102](#)

### Overview of GRE

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

GRE is described in RFC 2784 (obsoletes earlier RFCs 1701 and 1702). The switches support RFC 2784, but not completely. (For a list of limitations, see [“Configuration Limitations” on page 102.](#))

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

## GRE Tunneling

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

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

For details about GRE tunneling, see:

- [Encapsulation and De-Encapsulation on the Switch on page 100](#)
- [Number of Source and Destination Tunnels Allowed on a Switch on page 100](#)
- [Class of Service on GRE Tunnels on page 101](#)
- [Applying Firewall Filters to GRE Traffic on page 101](#)

---

### Encapsulation and De-Encapsulation on the Switch

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

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

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

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

---

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

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

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

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

### Class of Service on GRE Tunnels

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

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

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

### Applying Firewall Filters to GRE Traffic

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a switch. (For details, see [Firewall Filters for EX Series Switches Overview](#).) Because of the encapsulation and de-encapsulation performed by GRE, you are constrained as to where you can apply a firewall filter to filter tunneled packets and which header will be affected. [Table 21 on page 101](#) identifies these constraints.

**Table 21: Firewall Filter Application Points for Tunneled Packets**

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

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

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

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

## Configuration Limitations

Table 22 on page 102 lists features that are not supported with GRE.

**Table 22: Features Not Supported with GRE**

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

- Related Documentation**
- *Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)*
  - [Configuring Generic Routing Encapsulation Tunneling on page 102](#)
  - *Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch*

## Configuring Generic Routing Encapsulation Tunneling

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

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

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

This topic describes:

1. [Configuring a GRE Tunnel on page 103](#)

## Configuring a GRE Tunnel

To configure a GRE tunnel interface:

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

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



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

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

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

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

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

3. Specify the destination address:

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

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

### Related Documentation

- [Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly on page 103](#)
- [Understanding Generic Routing Encapsulation on page 99](#)
- [Configuring a Firewall Filter to De-encapsulate GRE Traffic on a QFX5100 Switch](#)

## Verifying That Generic Routing Encapsulation Tunneling Is Working Correctly

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

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

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

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

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

**Related Documentation**

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

## PART 4

# IP Directed Broadcast

- [Understanding IP Directed Broadcast on page 107](#)





## CHAPTER 5

# Understanding IP Directed Broadcast

- [Understanding IP Directed Broadcast on page 107](#)
- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 109](#)
- [Example: Configuring IP Directed Broadcast on page 110](#)

## Understanding IP Directed Broadcast

---

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

This topic covers:

- [IP Directed Broadcast Overview on page 107](#)
- [IP Directed Broadcast Implementation on page 108](#)
- [When to Enable IP Directed Broadcast on page 108](#)
- [When Not to Enable IP Directed Broadcast on page 108](#)

## IP Directed Broadcast Overview

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

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

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

## IP Directed Broadcast Implementation

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

By default, IP directed broadcast is disabled.

## When to Enable IP Directed Broadcast

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

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

## When Not to Enable IP Directed Broadcast

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

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

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

### Related Documentation

- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 109](#)
- [Example: Configuring IP Directed Broadcast on page 110](#)

## Configuring IP Directed Broadcast (CLI Procedure)

You can use IP directed broadcast on a switch to facilitate remote network management by sending broadcast packets to hosts on a specified subnet without broadcasting to the entire network. IP directed broadcast packets are broadcast on only the target subnet. The rest of the network treats IP directed broadcast packets as unicast packets and forwards them accordingly.

Before you begin to configure IP directed broadcast:

- Ensure that the subnet on which you want broadcast packets using IP direct broadcast is not directly connected to the Internet.
- Configure an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) for the subnet that will be enabled for IP direct broadcast. See *Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)*, *Configuring Routed VLAN Interfaces (CLI Procedure)*, or *Configuring VLANs*.



**NOTE:** We recommend that you do not enable IP directed broadcast on subnets that have a direct connection to the Internet because of increased exposure to denial-of-service (DoS) attacks.

To enable IP directed broadcast for a specified subnet:



**NOTE:** In a mixed Virtual Chassis, when you configure targeted broadcast, you can only configure one interface. Otherwise, targeted broadcast will not work.

1. Add the target subnet's logical interfaces to the VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```

2. Configure the Layer 3 interface on the VLAN that is the target of the IP directed broadcast packets:

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

3. Associate a Layer 3 interface with the VLAN:

```
[edit vlans]
user@switch# set v1 l3-interface irb.1
```

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

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

### Related Documentation

- [Example: Configuring IP Directed Broadcast on page 110](#)
- [Understanding IP Directed Broadcast on page 107](#)

## Example: Configuring IP Directed Broadcast

---

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

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

- [Requirements on page 110](#)
- [Overview and Topology on page 110](#)
- [Configuration on page 111](#)

### Requirements

This example uses the following software and hardware components:

- Junos OS Release 15.1 or later for QFX Series switches
- One PC
- One QFX Series switch

Before you configure IP directed broadcast for a subnet:

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

### Overview and Topology

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

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

In this topology (see [Figure 1 on page 111](#)), a host is connected to an interface on an EX Series switch to manage the clients in subnet **10.1.2.1/24**. When the switch receives a packet with the broadcast IP address of the target subnet as its destination address, it forwards the packet to the subnet's Layer 3 interface and broadcasts it to all the hosts within the subnet.

Figure 1: Topology for IP Directed Broadcast




Table 23 on page 111 shows the settings of the components in this example.

Table 23: Components of the IP Directed Broadcast Topology

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

Configuration

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

CLI Quick Configuration

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

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

Step-by-Step Procedure

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

1. Add logical interface **ge-0/0/0.0** to VLAN **v1**:  
[edit interfaces]  
user@switch# **set ge-0/0/0.0 family ethernet-switching vlan members v1**
2. Add logical interface **ge-0/0/1.0** to VLAN **v1**:  
[edit interfaces]

- ```

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

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


```

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

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


```

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

```
 5. Configure the IP address for the ingress VLAN:


```

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

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


```

[edit vlans]
user@switch# set v1 l3-interfacevlan.1
user@switch# set v0 l3-interface vlan.0

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


```

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

```

Results Check the results:

```

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

```

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

Related Documentation

- [Configuring IP Directed Broadcast \(CLI Procedure\) on page 109](#)

PART 5

Layer 3 Logical Interfaces

- [Understanding Layer 3 Logical Interfaces on page 117](#)

CHAPTER 6

Understanding Layer 3 Logical Interfaces

- [Understanding Layer 3 Logical Interfaces on page 117](#)
- [Configuring a Layer 3 Logical Interface on page 118](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 118](#)

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.

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

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

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

Related Documentation

- [Interfaces Overview on page 3](#)
- [Configuring a Layer 3 Logical Interface on page 118](#)
- *Junos OS Network Interfaces Library for Routing Devices*

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*.

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
```

Related Documentation

- [Understanding Layer 3 Logical Interfaces on page 117](#)
- [Verifying That Layer 3 Logical Interfaces Are Working on page 118](#)

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 | 1.1.1.1/24 | |
| ge-0/0/0.1 | up | up | inet | 2.1.1.1/24 | |
| ge-0/0/0.2 | up | up | inet | 3.1.1.1/24 | |
| ge-0/0/0.3 | up | up | inet | 4.1.1.1/24 | |
| ge-0/0/0.4 | up | up | inet | 5.1.1.1/24 | |
| ge-0/0/0.32767 | up | up | | | |

2. Use the **ping** command from a device on one subnet to an address on another subnet to determine if packets were transmitted correctly on the logical interface VLANs:

```
user@switch> ping ip-address
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
```

Meaning The output confirms that the logical interfaces have been created and the links are up.

Related Documentation

- [Configuring a Layer 3 Logical Interface on page 118](#)

PART 6

Local Link Bias

- [Understanding Local Link Bias on page 123](#)

CHAPTER 7

Understanding Local Link Bias

- [Understanding Local Link Bias on page 123](#)
- [Configuring Local Link Bias \(CLI Procedure\) on page 125](#)

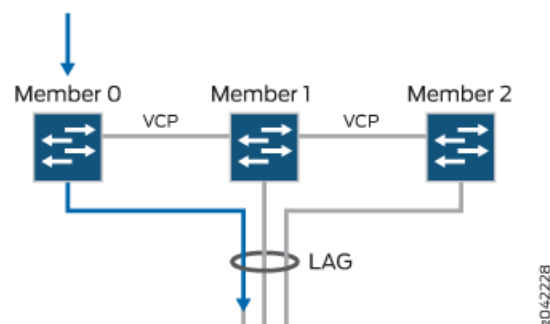
Understanding Local Link Bias



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

Local link bias conserves bandwidth on Virtual Chassis ports (VCPs) by using local links to forward unicast traffic exiting a Virtual Chassis or Virtual Chassis Fabric (VCF) that has a Link Aggregation group (LAG) bundle composed of member links on different member switches in the same Virtual Chassis or VCF. A local link is a member link in the LAG bundle that is on the member switch that received the traffic. Because traffic is received and forwarded on the same member switch when local link bias is enabled, no VCP bandwidth is consumed by traffic traversing the VCPs to exit the Virtual Chassis or VCF using a different member link in the LAG bundle. The traffic flow of traffic exiting a Virtual Chassis or VCF over a LAG bundle when local link bias is enabled is illustrated in [Figure 2 on page 123](#).


Figure 2: 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 3 on page 124](#).

Figure 3: Egress Traffic Flow without Local Link Bias



Local link bias is configured in a LAG bundle. 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.

Related Documentation

- [Configuring Local Link Bias \(CLI Procedure\) on page 125](#)

Configuring Local Link Bias (CLI Procedure)

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

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

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
```

Related Documentation

- [Understanding Local Link Bias on page 123](#)

PART 7

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

- [Understanding LAGs and LACP on page 129](#)

CHAPTER 8

Understanding LAGs and LACP

- Understanding Aggregated Ethernet Interfaces and LACP on page 129
- Configuring Aggregated Ethernet LACP on page 132
- Configuring Link Aggregation on page 133
- Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136
- Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 140
- Verifying the Status of a LAG Interface on page 145
- Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 145
- Troubleshooting an Aggregated Ethernet Interface on page 146

Understanding Aggregated Ethernet Interfaces and LACP

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

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



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



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

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



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



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

- [Link Aggregation Group on page 130](#)
- [Link Aggregation Control Protocol \(LACP\) on page 131](#)

Link Aggregation Group

To create a LAG:

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

Keep in mind these hardware and software guidelines:

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



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

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



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

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



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



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

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

Link Aggregation Control Protocol (LACP)

LACP is one method of bundling several physical interfaces to form one logical aggregated Ethernet interface. The LACP mode can be active or passive. The transmitting link is known as the *actor*, and the receiving link is known as the *partner*. If the actor and partner are both in passive mode, they do not exchange LACP packets, and the aggregated Ethernet links do not come up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode. You can configure Ethernet links to actively transmit protocol data units (PDUs), or you can configure the links to passively transmit them, sending out LACP

PDU only when they receive them from another link. You can configure both VLAN-tagged and untagged aggregated Ethernet interfaces without LACP enabled. LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

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

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

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

**Related
Documentation**

- [Configuring Link Aggregation on page 133](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136](#)
- [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
- [Verifying the Status of a LAG Interface on page 145](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

Configuring Aggregated Ethernet LACP

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

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

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



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

To configure LACP:

1. Enable the LACP mode:

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

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

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

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

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

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

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

Configuring Link Aggregation

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



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



NOTE: On QFX5100, 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.

1. [Creating an Aggregated Ethernet Interface on page 134](#)
2. [Configuring the VLAN Name and VLAN ID Number on page 134](#)
3. [Configuring Aggregated Ethernet LACP on page 135](#)

Creating an Aggregated Ethernet Interface

To create an aggregated Ethernet interface:

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

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

For example, to specify 5:

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

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



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

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

For example, to specify 5:

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

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

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

For example, to specify 10g:

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

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

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

Configuring the VLAN Name and VLAN ID Number



NOTE: VLANs are not supported on OCX Series switches.

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

For example, 100.

Configuring Aggregated Ethernet LACP

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

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

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



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

To configure LACP:

1. Enable the LACP mode:

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

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

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

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

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

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

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

Related Documentation

- [Understanding Interface Naming Conventions on page 5](#)
- [Configuring an FCoE LAG](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136](#)
- [Verifying the Status of a LAG Interface on page 145](#)
- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 145](#)
- [show lacp statistics interfaces \(View\) on page 433](#)

Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch

A QFX Series product allows you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. You can configure LAGs to connect a QFX Series product to other switches, like aggregation switches, servers, or routers. This example describes how to configure LAGs to connect a QFX3500, QFX3600, EX4600, QFX5100, and QFX10002 switch to an aggregation switch.

- [Requirements on page 136](#)
- [Overview and Topology on page 136](#)
- [Configuration on page 137](#)
- [Verification on page 139](#)
- [Troubleshooting on page 140](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX3500 and QFX3600 switches, Junos OS 13.2 or later for the QFX5100 and EX4600 switch, and Junos OS Release 15.1X53-D10 for QFX10002 switches.
- One QFX3500, QFX3600, EX4600, QFX5100, QFX10002 switch.

Overview and Topology

In this example, the switch has one LAG comprising two 10-Gigabit Ethernet interfaces. This LAG is configured in port mode trunk so that the switch and the VLAN to which it has been assigned can send and receive traffic.

Configuring the Ethernet interfaces as LAGs has the following advantages:

- If one physical port is lost for any reason (a cable is unplugged or a switch port fails), the logical port transparently continues to function over the remaining physical port.
- Link Aggregation Control Protocol (LACP) can optionally be configured for link monitoring and automatic addition and deletion of individual links without user intervention.



.....

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

.....

The topology used in this example consists of one switch with a LAG configured between two of its 10-Gigabit Ethernet interfaces. The switch is connected to an aggregation switch.




Table 24 on page 137 details the topology used in this configuration example.

Table 24: Components of the Topology for Configuring a LAG Between a QFX3500 Switch and Aggregation Switch

| Hostname | Base Hardware | Trunk Port |
|----------|---|--|
| switch | QFX3500, QFX3600, EX4600, QFX5100, or QFX10002 switch | ae0 is configured as a trunk port and combines the following two interfaces: xe-0/0/2 and xe-0/0/3 |

Configuration

To configure a LAG between two 10-Gigabit Ethernet interfaces:

CLI Quick Configuration

To quickly configure a LAG between two 10-Gigabit Ethernet interfaces on a switch, copy the following commands and paste them into the switch terminal window:



NOTE: If you are configuring a LAG using Enhanced Layer 2 software, use the **interface-mode** statement instead of the **port-mode** statement. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

```
[edit]
set chassis aggregated-devices ethernet device-count 1
set interfaces ae0 aggregated-ether-options minimum-links 1
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family ethernet-switching vlan members green
set interfaces xe-0/0/2 ether-options 802.ad ae0
set interfaces xe-0/0/3 ether-options 802.ad ae0
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 aggregated-ether-options lacp periodic fast
```

Step-by-Step Procedure

To configure a LAG between a QFX Series switch and an aggregation switch:

- Specify the number of LAGs to be created on the switch:

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 1
```
- Specify the number of links that need to be present for the **ae0** LAG interface to be up:

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

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed 10g
```
  4. Specify the members to be included within the aggregated Ethernet bundle:
 

```
[edit interfaces]
user@switch# set interfaces xe-0/0/2 ether-options 802.ad ae0
[edit interfaces]
user@switch# set interfaces xe-0/0/3 ether-options 802.ad ae0
```
  5. Assign a port mode of trunk to the **ae0** link:



**NOTE:** If you are configuring a LAG on the QFX5100 switch, use the **interface-mode** statement instead of the **port-mode** statement. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

- ```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching port-mode trunk
or
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching interface-mode trunk
```
6. Assign the LAG to a VLAN:


```
[edit interfaces]
user@switch# set ae0 unit 0 family ethernet-switching vlan members green vlan-id 200
```
 7. (Optional): Designate one side of the LAG as active for LACP:


```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp active
```
 8. (Optional): Designate the interval and speed at which the interfaces send LACP packets:


```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options lacp periodic fast
```

Results

Display the results of the configuration on a QFX3500 or QFX3600 switch:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 1;
    }
  }
}
green {
  vlan-id 200;
}
}
interfaces {
  ae0 {
```



```

aggregated-ether-options {
  link-speed 10g;
  minimum-links 1;
}
unit 0 {
  family ethernet-switching {
    port-mode trunk;
    vlan {
      members green;
    }
  }
}
xe-0/0/2 {
  ether-options {
    802.ad ae0;
  }
}
xe-0/0/3 {
  ether-options {
    802.ad ae0;
  }
}
}

```

Verification

To verify that switching is operational and one LAG has been created, perform these tasks:

- [Verifying That LAG ae0.0 Has Been Created on page 139](#)
- [Verifying That LAG ae0 Has Been Created on page 139](#)

Verifying That LAG ae0.0 Has Been Created

Purpose Verify that LAG **ae0.0** has been created on the switch.

Action `show interfaces ae0 terse`

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|------------|-------|--------|
| ae0 | up | up | | | |
| ae0.0 | up | up | eth-switch | | |

Meaning The output confirms that the **ae0.0** link is up and shows the **family** and IP address assigned to this link.

Verifying That LAG ae0 Has Been Created

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

Action `show interfaces ae0 terse`

| Interface | Admin | Link | Proto | Local | Remote |
|-----------|-------|------|------------|-------|--------|
| ae0 | up | down | | | |
| ae0.0 | up | down | eth-switch | | |

Meaning The output shows that the **ae0.0** link is down.

Troubleshooting

Troubleshooting a LAG That Is Down

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

Solution Check the following:

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

- Related Documentation**
- [Configuring Link Aggregation on page 133](#)
 - [Verifying the Status of a LAG Interface on page 145](#)
 - [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 145](#)
 - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 140](#)
 - [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
 - [show lacp statistics interfaces \(View\) on page 433](#)

Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch

QFX Series products allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your QFX Series product model. On a standalone switch, you can group up to 32 Ethernet interfaces to form a LAG. On a QFabric system, you can group up to 8 Ethernet interfaces to form a LAG. QFX Series products allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in [“Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch” on page 136](#):

- [Requirements on page 141](#)
- [Overview and Topology on page 141](#)
- [Configuring LACP for the LAG on the QFX Series on page 141](#)

- [Verification on page 142](#)
- [Troubleshooting on page 143](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 11.1 or later for the QFX3500 switch, Junos OS Release 12.1 or later for the QFX3600 switch, Junos OS Release 13.2 or later for the QFX5100 switch, and Junos OS Release 15.1X53-D10 or later for the QFX10002 switch.
- One QFX3500, QFX3600, QFX5100, QFX10002 switch.

Before you configure LACP, be sure you have:

- Configured the ports on the switches as trunk ports.
- Configured the LAG.

Overview and Topology

The topology in this example is exactly the same as the topology used in the [Configuring a LAG Between a QFX Switch and an Aggregation Switch](#) example. This example shows how to use LACP to enhance the LAG functionality.

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



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

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the **periodic** statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

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

Configuring LACP for the LAG on the QFX Series

To configure LACP for a QFX Series LAG, perform these tasks:

CLI Quick Configuration

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

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

Step-by-Step Procedure

To configure LACP for LAG ae0 :

1. Specify the aggregated Ethernet options for the LAG:

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

Results Display the results of the configuration:

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

Verification

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

- [Verifying the LACP Settings on page 142](#)
- [Verifying That the LACP Packets Are Being Exchanged on page 142](#)

Verifying the LACP Settings**Purpose** Verify that LACP has been set up correctly.**Action** Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@switch> show lacp interfaces xe-0/0/02
```

Aggregated interface: ae0

| | | | | | | | | | |
|----------------|---------------|----------------|-----|------|-----|-----|-----------|---------|----------|
| LACP state: | Role | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
| xe-0/0/2 | Actor | No | Yes | No | No | No | Yes | Fast | Active |
| xe-0/0/2 | Partner | No | Yes | No | No | No | Yes | Fast | Passive |
| LACP protocol: | Receive State | Transmit State | | | | | Mux State | | |
| xe-0/0/2 | Defaulted | Fast periodic | | | | | Detached | | |

Meaning The output indicates that LACP has been set up correctly and is active at one end.**Verifying That the LACP Packets Are Being Exchanged****Purpose** Verify that LACP packets are being exchanged.

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

```
user@switch> show interfaces ae0 statistics
```

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

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

Meaning The output here shows that the link is down and that no PDUs are being exchanged.

Troubleshooting

To troubleshoot a nonworking LACP link, perform these tasks:

- [Troubleshooting a Nonworking LACP Link on page 143](#)

[Troubleshooting a Nonworking LACP Link](#)

Problem The LACP link is not working.

Solution Check the following:

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

Related Documentation

- [Configuring Link Aggregation on page 133](#)
- [Verifying the Status of a LAG Interface on page 145](#)

- [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 145](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136](#)
- [Example: Configuring an FCoE LAG on a Redundant Server Node Group](#)
- [show lacp statistics interfaces \(View\) on page 433](#)

Verifying the Status of a LAG Interface

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

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

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

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

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

- Related Documentation**
- [Configuring Link Aggregation on page 133](#)
 - [Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets on page 145](#)
 - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136](#)
 - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 140](#)
 - [show lacp statistics interfaces \(View\) on page 433](#)

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 145](#)
2. [Verifying That LACP Packets Are Being Exchanged on page 146](#)

Verifying the LACP Setup

Purpose Verify that the LACP has been set up correctly.

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

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

```
Aggregated interface: ae0
```

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

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

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

Verifying That LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged between interfaces.

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

`show lacp statistics interfaces ae0`

Aggregated interface: ae0

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

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

- Related Documentation**
- [Configuring Link Aggregation on page 133](#)
 - [Verifying the Status of a LAG Interface on page 145](#)
 - [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136](#)
 - [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 140](#)
 - [show lacp statistics interfaces \(View\) on page 433](#)

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 145](#)
- [Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136](#)

PART 8

Redundant Trunk Groups

- [Understanding Redundant Trunk Groups on page 151](#)

CHAPTER 9

Understanding Redundant Trunk Groups

- [Understanding Redundant Trunk Links on page 151](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 153](#)

Understanding Redundant Trunk Links

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.

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 4 on page 152](#), 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 4 on page 152](#) 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 4: Redundant Trunk Group, Link 1 Active





Figure 5 on page 152 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 5: 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 1 and Switch 2.

**Related
Documentation**

- [Example: Configuring Redundant Trunk Links for Faster Recovery](#)
- [Example: Configuring Redundant Trunk Links for Faster Recovery on page 153](#)

Example: Configuring Redundant Trunk Links for Faster Recovery



NOTE: This example uses Junos OS for EX Series switches or QFX Series with support for the Enhanced Layer 2 Software (ELS) configuration style. If your EX Series switch runs software that does not support ELS, see *Example: Configuring Redundant Trunk Links for Faster Recovery*. For ELS details, see *Getting Started with Enhanced Layer 2 Software*.

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 153](#)
- [Overview and Topology on page 154](#)
- [Disabling RSTP on Switches 1 and 2 on page 156](#)
- [Configuring Redundant Trunk Links on Switch 3 on page 156](#)
- [Verification on page 157](#)

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 6 on page 155](#)).

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 6 on page 155 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 25 on page 156 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 6: Topology for Configuring the Redundant Trunk Links




Table 25: 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 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 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 157](#)

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.

Related Documentation

- [Understanding Redundant Trunk Links on page 151](#)

PART 9

Resilient Hashing

- [Understanding Resilient Hashing on page 161](#)

CHAPTER 10

Understanding Resilient Hashing

- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 161](#)
- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 167](#)
- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 169](#)
- [Configuring Resilient Hashing for Trunk/ECMP Groups on page 171](#)

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

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

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

This topic contains the following sections:

- [Understanding the Hashing Algorithm on page 161](#)
- [IP \(IPv4 and IPv6\) on page 162](#)
- [MPLS on page 164](#)
- [MAC-in-MAC Packet Hashing on page 165](#)
- [Layer 2 Header Hashing on page 166](#)

Understanding the Hashing Algorithm

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

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

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

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

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

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

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

Note the following points regarding the hashing algorithm:

- The fields selected for hashing are based on the packet type only. The fields are not based on any other parameters, including forwarding decision (bridged or routed) or egress LAG bundle configuration (Layer 2 or Layer 3).
- The same fields are used for hashing unicast and multicast packets. Unicast and multicast packets are, however, hashed differently.
- The same fields are used by the hashing algorithm to hash ECMP and LAG traffic, but the hashing algorithm hashes ECMP and LAG traffic differently. LAG traffic uses a trunk hash while ECMP uses ECMP hashing. Both LAG and ECMP use the same RTAG7 seed but use different offsets of that 128B seed to avoid polarization. The initial config of the HASH function to use the trunk and ECMP offset are set at the PFE Init time. The different hashing ensures that traffic is not polarized when a LAG bundle is part of the ECMP next-hop path.
- The same fields are used for hashing regardless of whether the switch is or is not participating in a mixed or non-mixed Virtual Chassis or Virtual Chassis Fabric (VCF).

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

IP (IPv4 and IPv6)

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

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

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

Table 26 on page 163 displays the IPv4 and IPv6 payload fields that are used by the hashing algorithm, by default.

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

Table 26: IPv4 and IPv6 Hashing Fields

| Fields | EX4300 | | QFX5100 | | QFX5110 | | QFX5200 | |
|-------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| | LAG | ECMP | LAG | ECMP | LAG | ECMP | LAG | ECMP |
| Source MAC | X | X | X | X | X | X | X | X |
| Destination MAC | X | X | X | X | X | X | X | X |
| EtherType | X | X | X | X | X | X | X | X |
| VLAN ID | X (configurable) | X (configurable) | X (configurable) | X (configurable) | X (configurable) | X (configurable) | X (configurable) | X (configurable) |
| Source IP or IPv6 | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | X (configurable) | X (configurable) |
| Destination IP or IPv6 | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | X (configurable) | X (configurable) |
| Protocol (IPv4 only) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | X (configurable) | X (configurable) |
| Next header (IPv6 only) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | X (configurable) | X (configurable) |
| Layer 4 Source Port | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | ✓ (configurable) | X (configurable) | X (configurable) |

Table 26: IPv4 and IPv6 Hashing Fields (*continued*)

| Fields | EX4300 | | QFX5100 | | QFX5110 | | QFX5200 | |
|--------------------------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------------|
| Layer 4 Destination Port | ✓ (config) | ✓ (config) | ✓ (config) | ✓ (config) | ✓ (config) | ✓ (config) | X (config) | X (configurable) |
| IPv6 Flow label (IPv6 only) | X | X | X | X | X | X | X | X |

MPLS

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

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

Table 27 on page 164 displays the MPLS payload fields that are used by the hashing algorithm, by default:

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

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

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

Table 27: MPLS Hashing Fields

| Field | EX4300 | QFX5100 | QFX5110 | QFX5200 |
|-----------------|--------|---------|---------|---------|
| Source MAC | X | X | X | X |
| Destination MAC | X | X | X | X |
| EtherType | X | X | X | X |
| VLAN ID | X | X | X | X |
| Source IP | ✓ | ✓ | ✓ | ✓ |
| Destination IP | ✓ | ✓ | ✓ | ✓ |

Table 27: MPLS Hashing Fields (*continued*)

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

MAC-in-MAC Packet Hashing

Packets using the MAC-in-MAC EtherType are hashed by the hashing algorithm using the Layer 2 payload source MAC, Layer 2 payload destination MAC, and Layer 2 payload EtherType fields. See [Table 28 on page 165](#).

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

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

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

Table 28: MAC-in-MAC Hashing Fields

| Field | EX4300 | QFX5100 | QFX5110 | QFX5200 |
|----------------------------|--------|---------|---------|---------|
| Layer 2 Payload Source MAC | ✓ | ✓ | ✓ | ✓ |

Table 28: MAC-in-MAC Hashing Fields (*continued*)

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

Layer 2 Header Hashing

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

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

Table 29: Layer 2 Header Hashing Fields

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

Related Documentation

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 169](#)

Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups

You use resilient hashing to minimize flow remapping across members of a trunk/ECMP group in a load-balanced system. You can configure resilient hashing in link aggregation groups (LAGs) and in equal cost multipath (ECMP) groups.



NOTE: Resilient hashing in LAGs is not supported on QFX10002 and QFX 10008 Switches.

- [Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing on page 167](#)
- [Limitations and Caveats for Resilient Hashing on page 168](#)
- [Resilient Hashing on LAGs on page 168](#)
- [Resilient Hashing on ECMP on page 169](#)

Why You Might Want to Use Resilient Hashing and How It Works with Static Hashing

Resilient hashing works in conjunction with the default static hashing algorithm. When members are added to or deleted from a trunk/ECMP group, the static hashing algorithm might remap destination paths. Resilient hashing distributes traffic across all members of a group by tracking the flow's member utilization. When a flow is affected by a member change, the Packet Forwarding Engine rebalances the flow by reprogramming the flow set table.

Resilient hashing thus provides the following benefits:

- Minimizes traffic-distribution imbalances among members of a trunk/ECMP group when members are added to or deleted from the group.
- Minimizes the impact on flows bound to unaffected members when a new member is added or an existing member is deleted from the group.

In normal hash-based load balancing, with the static hashing algorithm used alone, flows are assigned to members through the mathematical mod (%) operation. Any increase or decrease in the number of group members results in a complete remapping of flows to member IDs, as shown in the following example:

- Member ID = Hash (key) mod (number of members in group)
- Example:
 - Hash (key) = 10
 - $10 \bmod 5 = 0$ (member with ID 0 is selected for flow)
 - $10 \bmod 4 = 2$ (member with ID 2 is selected for the same flow when the number of members is decreased by 1)

Resilient hashing minimizes the destination path remapping when a member in the trunk/ECMP group is added or deleted.

When the flow is affected by a member change in the group, resilient hashing rebalances the flow by reprogramming the flow set table.

Table 30: Destination Path Results for Static Hashing and for Resilient Hashing When Members Are Added to or Deleted from Trunk Groups

| Trunk Group Size | Normal (Static) Hashing Result | Resilient Hashing Result | Notes |
|------------------|--|---|---|
| 4 | Hash(10) % 4 = 2 Flow is assigned to member ID 2. | Flow is assigned to one of four group members based on flow set table entries. | Original trunk/ECMP group size is 4. |
| 3 | Hash(10) % 3 = 1 Flow is assigned to member ID 1. | Flow is assigned to same member as in the previous case. | Delete one member from original trunk/ECMP group. Trunk/ECMP group size is 3. |
| 5 | Hash(10) % 5 = 0 Flow is assigned to member ID 0. | There is minimal redistribution of flows from other members to this newly added member. | Add one member to original trunk group. Trunk/ECMP group size is 5. |

Limitations and Caveats for Resilient Hashing

Notice the following limitation and caveats for the resilient hashing feature:

- Resilient hashing applies only to unicast traffic.
- Resilient hashing supports a maximum of 1024 trunk groups, with each group having a maximum of 256 members.
- Resilient hashing does not guarantee that traffic distribution is even across all group members—it depends on the traffic pattern and on the organization of the resilient hashing flow set table in hardware. Resilient hashing *minimizes* remapping of flows to destination links when members are added to or deleted from the group.
- If resilient hashing is enabled on a trunk group or ECMP group and if **set forwarding-options enhanced-hash-key** with one of the options **hash-mode**, **inet**, **inet6**, or **layer2** is used, some flows might change destination links, because the new hash parameters might generate new hash indexes for the flows, and hence the new destination links.
- Resilient hashing is not supported on Virtual Chassis port (VCP) links.

Resilient Hashing on LAGs

A LAG combines Ethernet interfaces (members) to form a logical point-to-point link that increases bandwidth, provides reliability, and allows load balancing. Resilient hashing minimizes destination remapping behavior when a new member is added or deleted from the LAG.

A resilient hashing configuration on LAGs is per-aggregated-Ethernet-interface–based.

Resilient Hashing on ECMP

An ECMP group for a route contains multiple next-hop equal cost addresses for the same destination in the routing table. (Routes of equal cost have the same preference and metric values.)

Junos OS uses the static hashing algorithm to choose one of the next-hop addresses in the ECMP group to install in the forwarding table. Resilient hashing enhances ECMPs by minimizing destination remapping behavior when a new member is added or deleted from the ECMP group.

A resilient hashing configuration on ECMP is global—it applies to all ECMP groups.

Related Documentation

- [Configuring Resilient Hashing for Trunk/ECMP Groups on page 171](#)

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

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

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

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



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

- [Configuring the Hashing Algorithm to Use Fields in the Layer 2 Header for Hashing on page 169](#)
- [Configuring the Hashing Algorithm to Use Fields in the IP Payload for Hashing on page 170](#)
- [Configuring the Hashing Algorithm to Use Fields in the IPv6 Payload for Hashing on page 170](#)

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

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

1. Configure the hash mode to Layer 2 header:

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

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

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

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

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

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

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

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

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

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

1. Configure the hash mode to Layer 2 payload:

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

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

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

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

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

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

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

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

1. Configure the hash mode to Layer 2 payload:

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

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

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


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

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

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

Related Documentation

- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 161](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic \(QFX 10002 and QFX 10008 Switches\)](#)
- [Understanding Aggregated Ethernet Interfaces and LACP](#)

Configuring Resilient Hashing for Trunk/ECMP Groups

You use resilient hashing to minimize flow remapping across members of a trunk/ECMP group in a load-balanced system. You can configure resilient hashing in link aggregation groups (LAGs) and in equal cost multipath (ECMP) sets.

This topic includes:

1. [Configuring Resilient Hashing on LAGs on page 171](#)
2. [Configuring Resilient Hashing on ECMP Groups on page 171](#)

Configuring Resilient Hashing on LAGs



NOTE: Configuring resilient hashing on LAGs is not supported on QFX10002 and QFX10008 switches.

To enable resilient hashing for a LAG:

- Configure resilient hashing on the aggregated Ethernet interface:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options resilient-hash
```

Configuring Resilient Hashing on ECMP Groups

To enable resilient hashing for ECMP groups:

- Configure resilient hashing for ECMP:

```
[edit forwarding-options]
user@switch# set enhanced-hash-key ecmp-resilient-hash
```

Related Documentation

- [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups on page 167](#)

PART 10

Uplink Failure Detection

- [Understanding Uplink Failure Detection on page 175](#)

CHAPTER 11

Understanding Uplink Failure Detection

- [Overview of Uplink Failure Detection on page 175](#)
- [Configuring Interfaces for Uplink Failure Detection on page 177](#)
- [Example: Configuring Interfaces for Uplink Failure Detection on page 178](#)
- [Verifying That Uplink Failure Detection Is Working Correctly on page 182](#)

Overview of Uplink Failure Detection

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate this information to the downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all of the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects disabled downlink interfaces, it switches over to the secondary link to help ensure that the traffic of the failed link is not dropped.

This topic describes:


- [Uplink Failure Detection Configuration on page 175](#)
- [Failure Detection Pair on page 176](#)

Uplink Failure Detection Configuration

Uplink failure detection allows switches to monitor uplink interfaces to spot link failures. When a switch detects a link failure, it automatically disables the downlink interfaces bound to the uplink interface. A server that is connected to the disabled downlink interface triggers a network adapter failover to a secondary link to avoid any traffic loss.

[Figure 7 on page 176](#) illustrates a typical setup for uplink failure detection.

Figure 7: Uplink Failure Detection Configuration on Switches



For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down all downlink interfaces within that group. If any uplink interface returns to service, then the switch brings all downlink interfaces in that group back to service.

The switch can monitor both physical interface links and logical interface links for uplink failures, but you must put the two types of interfaces into separate groups.



NOTE: For logical interfaces, the server must send keepalives between the switch and the server to detect failure of logical links.

Failure Detection Pair

Uplink failure detection requires that you create pairs of uplink and downlink interfaces in a group. Each pair includes one of each of the following:

- A link-to-monitor interface—The link-to-monitor interfaces specify the uplinks the switch monitors. You can configure a maximum of 48 uplink interfaces as link-to-monitor interfaces for a group.
- A link-to-disable interface—The link-to-disable interfaces specify the downlinks the switch disables when the switch detects an uplink failure. You can configure a maximum of 48 downlinks to disable in the group.

The link-to-disable interfaces are bound to the link-to-monitor interfaces within the group. When a link-to-monitor interface returns to service, the switch automatically enables all link-to-disable interfaces in the group.

- Related Documentation**
- [Configuring Interfaces for Uplink Failure Detection on page 177](#)
 - [Example: Configuring Interfaces for Uplink Failure Detection on page 178](#)

Configuring Interfaces for Uplink Failure Detection

You can configure uplink failure detection to help ensure balanced traffic flow. Using this feature, switches can monitor and detect link failure on uplink interfaces and can propagate the failure information to downlink interfaces, so that servers connected to those downlinks can switch over to secondary interfaces.

Follow these configuration guidelines:

- Configure an interface in only one group.
- Configure a maximum of 48 groups for each switch.
- Configure a maximum of 48 uplinks to monitor and a maximum of 48 downlinks to disable in each group.
- Configure physical links and logical links in separate groups.

To configure uplink failure detection on a switch:

1. Specify a name for an uplink failure detection group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name
```

2. Add an uplink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-monitor interface-name
```

3. Repeat Step 2 for each uplink interface you add to the group.

4. Add a downlink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-disable interface-name
```

5. Repeat Step 4 for each downlink interface you add to the group.



NOTE: After you have configured an uplink failure detection group, use the `show uplink-failure-detection group (Uplink Failure Detection) group-name` command to verify that all interfaces in the group are up. If the interfaces are down, uplink failure detection does not work.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 175](#)
 - [Example: Configuring Interfaces for Uplink Failure Detection on page 178](#)
 - [Verifying That Uplink Failure Detection Is Working Correctly on page 182](#)

Example: Configuring Interfaces for Uplink Failure Detection

Uplink failure detection allows a switch to detect link failure on uplink interfaces and to propagate the failure information to the downlink interfaces. All of the network interface cards (NICs) on a server are configured as being either the primary link or the secondary link and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link to ensure that the traffic on the failed link is not dropped.

This example describes:

- [Requirements on page 178](#)
- [Overview and Topology on page 178](#)
- [Configuring Uplink Failure Detection on Both Switches on page 179](#)
- [Verification on page 180](#)

Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.1 or later for the QFX Series
- Two QFX3500 switches
- Two aggregation switches
- One dual-homed server

Overview and Topology

The topology in this example illustrates how to configure uplink failure detection on Switch A and Switch B. Switch A and Switch B are both configured with a link-to-monitor interface (the uplink interface to the aggregation switch) and a link-to-disable interface (the downlink interface to the server). For simplicity, only one group of link-to-monitor interfaces and link-to-disable interfaces is configured for each switch. The server is dual-homed to both Switch A and Switch B. In this scenario, if the link-to-monitor interface to Switch A is disabled, the server uses the link-to-monitor interface to Switch B instead.



NOTE: This example does not describe how to configure the dual-homed server or the aggregation switches. Please refer to the documentation for each of these devices for more information.

[Figure 7 on page 176](#) illustrates a typical setup for uplink failure detection.

Figure 8: Uplink Failure Detection Configuration on Switches




Table 31 on page 179 lists uplink failure settings for each QFX3500 switch.

Table 31: Settings for Uplink Failure Protection Example

| Switch A | Switch B |
|--|--|
| <ul style="list-style-type: none"> Group name: Group1 Link-to-monitor interface: xe-0/0/0 Link-to-disable interface: xe-0/0/1 | <ul style="list-style-type: none"> Group name: Group2 Link-to-monitor interface: xe-0/0/0 Link-to-disable interface: xe-0/0/1 |

Configuring Uplink Failure Detection on Both Switches

To configure uplink failure detection on both switches, perform these tasks:

CLI Quick Configuration To quickly configure uplink failure protection on Switch A and Switch B, copy the following commands and paste them into the switch terminal window:

```
[edit protocols]
set uplink-failure-detection group group1
set uplink-failure-detection group group2
set uplink-failure-detection group group1 link-to-monitor xe-0/0/0
set uplink-failure-detection group group2 link-to-monitor xe-0/0/0
set uplink-failure-detection group group1 link-to-disable xe-0/0/1
set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

Step-by-Step Procedure To configure uplink failure protection on both switches:

- Specify a name for the uplink failure detection group on Switch A:


```
[edit protocols]
user@switch# set uplink-failure-detection group group1
```
- Add an uplink interface to the group on Switch A:


```
[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-monitor xe-0/0/0
```

3. Add a downlink interface to the group on Switch A:

```
[edit protocols]
user@switch# set uplink-failure-detection group group1 link-to-disable xe-0/0/1
```
4. Specify a name for the uplink failure detection group on Switch B:

```
[edit protocols]
user@switch# set uplink-failure-detection group group2
```
5. Add an uplink interface to the group on Switch B:

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-monitor xe-0/0/0
```
6. Add a downlink interface to the group on Switch B:

```
[edit protocols]
user@switch# set uplink-failure-detection group group2 link-to-disable xe-0/0/1
```

Results Display the results of the configuration:

```
uplink-failure-detection {
  group {
    group1 {
      link-to-monitor {
        xe-0/0/0;
      }
      link-to-disable {
        xe-0/0/1;
      }
    }
    group2 {
      link-to-monitor {
        xe-0/0/0;
      }
      link-to-disable {
        xe-0/0/1;
      }
    }
  }
}
```

Verification

To verify that uplink failure detection is working correctly, perform the following tasks on Switch A and Switch B:

- [Verifying That Uplink Failure Detection is Working Correctly on page 180](#)

Verifying That Uplink Failure Detection is Working Correctly

Purpose Verify that the switch disables the downlink interface when it detects an uplink failure.

- Action** 1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : xe-0/0/0*
Downlink             : xe-0/0/1*
Failure Action       : Inactive
```



NOTE: The asterisk (*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : xe-0/0/0
Downlink             : xe-0/0/1
Failure Action       : Active
```

Meaning The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 175](#)
 - [Configuring Interfaces for Uplink Failure Detection on page 177](#)

Verifying That Uplink Failure Detection Is Working Correctly

Purpose Verify that the switch disables the downlink interface when it detects an uplink failure.

Action 1. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : xe-0/0/0*
Downlink             : xe-0/0/1*
Failure Action       : Inactive
```



NOTE: The asterisk (*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface xe-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink failure detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : xe-0/0/0
Downlink             : xe-0/0/1
Failure Action       : Active
```

Meaning The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down (there are no asterisks after the interface name) and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Overview of Uplink Failure Detection on page 175](#)
 - [Configuring Interfaces for Uplink Failure Detection on page 177](#)
 - [Example: Configuring Interfaces for Uplink Failure Detection on page 178](#)

PART 11

Configuration Statements and Operational Commands

- [Ethernet OAM Link Fault Management Configuration Statements on page 185](#)
- [GRE Configuration Statements on page 207](#)
- [Interfaces Configuration Statements on page 209](#)
- [IP Directed Broadcast Configuration Statement on page 275](#)
- [LAGs and LACP Configuration Statements on page 277](#)
- [Local Link Bias Configuration Statements on page 289](#)
- [Redundant Trunk Groups Configuration Statements on page 291](#)
- [Resilient Hashing Configuration Statements on page 297](#)
- [Uplink Failure Detection Configuration Statements on page 309](#)
- [Interfaces Operational Commands on page 313](#)
- [LAGs and LACP Operational Commands on page 427](#)
- [Redundant Trunk Group Operational Commands on page 435](#)
- [Resilient Hashing Operational Commands on page 439](#)
- [Uplink Failure Detection Operational Commands on page 445](#)

CHAPTER 12

Ethernet OAM Link Fault Management Configuration Statements

- [action \(OAM LFM\) on page 186](#)
- [action-profile on page 187](#)
- [allow-remote-loopback on page 188](#)
- [ethernet \(OAM LFM\) on page 189](#)
- [event-thresholds on page 191](#)
- [event \(OAM LFM\) on page 192](#)
- [frame-error on page 192](#)
- [frame-period on page 193](#)
- [frame-period-summary on page 193](#)
- [oam on page 194](#)
- [interface \(OAM LFM\) on page 196](#)
- [link-adjacency-loss on page 197](#)
- [link-discovery on page 197](#)
- [link-down on page 198](#)
- [link-event-rate on page 198](#)
- [link-fault-management on page 199](#)
- [negotiation-options on page 200](#)
- [no-allow-link-events on page 200](#)
- [pdu-interval on page 201](#)
- [pdu-threshold on page 201](#)
- [remote-loopback on page 202](#)
- [symbol-period on page 202](#)
- [syslog \(OAM LFM\) on page 203](#)
- [traceoptions \(OAM LFM\) on page 204](#)

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.</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 (CLI Procedure) on page 84 |

action-profile

| | |
|---------------------------------|---|
| Syntax | <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 | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.</p> <p>The remaining statements are explained separately.</p> |
| Options | <i>profile-name</i> —Name of the action profile. |
| 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"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

allow-remote-loopback

| | |
|---------------------------------|--|
| Syntax | allow-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 | 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. |
| 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">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

ethernet (OAM LFM)

```

Syntax  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 {
      delay delay-value;
      delay-variation delay-variation-value;
    }
    cycle-time cycle-time-value;
    iteration-period iteration-period-value;
    measurement-type two-way-delay;
    passive;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
}
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;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}

```

```

    }
  }
}

```

| | |
|---------------------------------|--|
| Hierarchy Level | [edit protocols oam] |
| 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) 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. |
| 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"> • Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86 • Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 • Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) |

event-thresholds

| | |
|---------------------------------|---|
| Syntax | <pre> event-thresholds { frame-error count; frame-period count; frame-period-summary count; symbol-period count; } </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 | Configure threshold limit values for link events in periodic OAM PDUs. The remaining statements are explained separately. |
| Required Privilege Level | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

event (OAM LFM)

| | |
|---------------------------------|---|
| Syntax | <pre>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 action-profile <i>profile-name</i>] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | <p>Configure link events in an action profile for Ethernet OAM link fault management (LFM).</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

frame-error

| | |
|---------------------------------|--|
| Syntax | <pre>frame-error <i>count</i>;</pre> |
| 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 (CLI Procedure) on page 84 |

frame-period

| | |
|---------------------------------|--|
| Syntax | <code>frame-period 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 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 (CLI Procedure) on page 84 |

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 (CLI Procedure) on page 84 |

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.lag Operation, Administration, and Management (OAM) support on the switches.

The remaining statements are explained separately.

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

- Related Documentation**
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86](#)
 - [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches](#)
 - [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 84](#)
 - [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\)](#)

interface (OAM LFM)

Syntax

```
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 link-fault-management](#)]

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.

Options *interface-name*—Name of the interface to be enabled for IEEE 802.3ah 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**
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86](#)
 - [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 84](#)

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 | 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"> • Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86 • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

link-discovery

| | |
|---------------------------------|--|
| Syntax | link-discovery (active passive); |
| 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 | 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 | 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"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

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 | 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">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

link-event-rate

| | |
|---------------------------------|--|
| Syntax | link-event-rate { frame-error <i>count</i> ; frame-period <i>count</i> ; frame-period-summary <i>count</i> ; symbol-period <i>count</i> ; } |
| 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. |
| 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">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

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.

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 on EX Series Switches on page 86](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 84](#)

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.</p> |
| 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">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

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 | 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">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

pdu-interval

| | |
|---------------------------------|---|
| Syntax | <code>pdu-interval <i>interval</i>;</code> |
| 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 periodic OAM PDU sending interval for fault detection. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. |
| Options | <i>interval</i> —Periodic OAM PDU sending interval. Range: 400 through 1000 milliseconds Default: 1000 milliseconds |
| 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"> • Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86 • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

pdu-threshold

| | |
|---------------------------------|--|
| Syntax | <code>pdu-threshold <i>threshold-value</i>;</code> |
| 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 | 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. |
| Options | <i>threshold-value</i> —Number of PDUs missed before declaring the peer lost. Range: 3 through 10 PDUs Default: 3 PDUs |
| 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"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

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">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

symbol-period

| | |
|---------------------------------|---|
| Syntax | symbol-period <i>count</i> ; |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management action-profile <i>profile-name</i> ; 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 | Configure the threshold for sending symbol period events or taking the action specified in the action profile. 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. |
| Options | <i>count</i> —Threshold count in seconds for symbol period events. Range: 1 through 100 seconds |
| 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">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

syslog (OAM LFM)

| | |
|---------------------------------|---|
| Syntax | syslog; |
| Hierarchy Level | [edit protocols oam ethernet link-fault-management action-profile <i>profile-name</i> ; action] |
| Release Information | Statement introduced in Junos OS Release 9.4 for EX Series switches. |
| Description | Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event. |
| 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">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |

traceoptions (OAM LFM)

Syntax traceoptions {
 file *filename* <files *number*> <match *regex*> <size *size*> <world-readable |
 no-world-readable>;
 flag *flag* ;
 no-remote-trace;
 }

Release Information Statement introduced in JUNOS Release 10.2 for EX Series switches.

Description Configure tracing options the link fault management.

Options 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 */var/log*.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum *xk* to specify KB, *xm* to specify MB, or *xg* 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 **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **action-profile**—Trace action profile invocation events.
- **all**—Trace all events.
- **configuration**—Trace configuration events.
- **protocol**—Trace protocol processing events.
- **routing socket**—Trace routing socket events.

match—(Optional) Refine the output to log only those lines that match the given regular expression.

no-world-readable—(Optional) Restrict file access to the user who created the file.

no-remote-trace—(Optional) Disable the remote trace.

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 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.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

Default: If you do not include this option, tracing output is appended to an existing trace file.

world-readable—(Optional) Enable unrestricted file access.

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

| | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 86• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 84 |
|------------------------------|--|

GRE Configuration Statements

- [destination \(Tunnels\) on page 207](#)
- [tunnel on page 208](#)
- [tunnel-port on page 208](#)

destination (Tunnels)

| | |
|---------------------------------|---|
| Syntax | <code>destination address;</code> |
| Hierarchy Level | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection. |
| Options | <i>address</i> —Address of the remote side of the connection. |
| 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 Interface Address on page 51 • Configuring Generic Routing Encapsulation Tunneling (CLI Procedure) • Junos OS Services Interfaces Library for Routing Devices • point-to-point |

tunnel

| | |
|---------------------------------|--|
| Syntax | <pre>tunnel { destination <i>destination-address</i>; source <i>source-address</i>; ttl <i>number</i>; (not supported on QFX and OCX Series switches) }</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] |
| Release Information | Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | <p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic.</p> <p>The remaining statements are explained separately.</p> |
| 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 Generic Routing Encapsulation Tunneling (CLI Procedure)</i> |

tunnel-port

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

Interfaces Configuration Statements

- [address on page 211](#)
- [alarm \(chassis\) on page 215](#)
- [auto-negotiation on page 216](#)
- [ccc on page 217](#)
- [channel-speed on page 218](#)
- [configured-flow-control on page 219](#)
- [craft-lockout on page 220](#)
- [description on page 221](#)
- [ethernet \(Alarm\) on page 222](#)
- [ethernet-switching on page 223](#)
- [ether-options on page 224](#)
- [eui-64 on page 225](#)
- [family on page 226](#)
- [fibre-channel \(Alarm\) on page 230](#)
- [filter on page 231](#)
- [flow-control on page 233](#)
- [fpc on page 234](#)
- [gratuitous-arp-reply on page 235](#)
- [hold-time \(Physical Interface\) on page 236](#)
- [irb \(Interfaces\) on page 238](#)
- [inet \(interfaces\) on page 241](#)
- [inet6 \(interfaces\) on page 242](#)
- [interface-mode on page 243](#)
- [interface-range on page 245](#)
- [interfaces on page 247](#)
- [link-down on page 254](#)
- [link-mode on page 255](#)
- [link-speed on page 256](#)

- [loopback \(Aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet\) on page 257](#)
- [mac on page 257](#)
- [management-ethernet \(Alarm\) on page 258](#)
- [member on page 259](#)
- [member-range on page 260](#)
- [mtu on page 261](#)
- [no-gratuitous-arp-request on page 262](#)
- [pic on page 263](#)
- [rx-buffers on page 264](#)
- [source on page 265](#)
- [speed on page 266](#)
- [targeted-broadcast on page 267](#)
- [traceoptions \(Individual Interfaces\) on page 268](#)
- [tx-buffers on page 269](#)
- [unit on page 271](#)
- [vlan-id on page 272](#)
- [vlan-tagging on page 273](#)

address

```

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

```

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

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the interface address.



NOTE: The vrrp High Availability functionality is not available on the QFX Series.

Options *address*—Address of the interface.

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



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

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

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

For more information on configuring the same address on multiple interfaces, see [“Configuring the Interface Address” on page 51](#).

The remaining statements are explained separately.



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


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

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

alarm (chassis)

| | |
|---------------------------------|--|
| Syntax | <pre>alarm { interface-type { alarm-name (ignore red yellow); } }</pre> |
| Hierarchy Level | [edit chassis], [edit chassis interconnect-device <i>name</i>], [edit chassis node-group <i>name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.2 for the ACX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| 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 <i>alarm-name</i> lines.</p> |
| Options | <p><i>alarm-name</i>—Alarm condition. For a list of conditions, see <i>System-Wide Alarms and Alarms for Each Interface Type</i>.</p> <p><i>ignore</i>—The specified alarm condition does not set off any alarm.</p> <p><i>interface-type</i>—Type of interface on which you are configuring the alarm: 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 | 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>Chassis Conditions That Trigger Alarms</i> <i>Chassis Alarm Messages on a QFX3500 Device</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. |
| 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. |
| | <div> 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.</div> |
| | <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 266• Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i>• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

CCC

| | |
|----------------------------|---|
| Syntax | ccc; |
| Hierarchy Level | [edit interfaces <i>ge-fpc/slot/ port</i> unit <i>logical-unit-number</i> family] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Configure the logical interface as a circuit cross-connect (CCC). |




NOTE: On QFX10002 switches, circuit cross-connects are not supported on aggregated Ethernet interfaces.

| | |
|---------------------------------|---|
| 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. |

channel-speed

| | |
|---------------------------------|---|
| Syntax | channel-speed (10g 25g 50g; 100g disable-auto-speed-detection) ; |
| Hierarchy Level | [edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> (port <i>port-number</i> port-range <i>port-range-low</i> <i>port-range-high</i>)] |
| 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">• Channelizing Interfaces on page 63• Channelizing Interfaces on QFX5200 Switches on page 75 |

configured-flow-control

| | |
|---------------------------------|--|
| Syntax | <pre>configured-flow-control { rx-buffers (on off); tx-buffers (on off); }</pre> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> 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 statements are explained separately. |
| 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 233 • <i>Configuring CoS Asymmetric Ethernet PAUSE Flow Control</i> • <i>Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control</i> • <i>Understanding CoS Flow Control (Ethernet PAUSE and PFC)</i> |

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

| | |
|---------------------------------|---|
| Syntax | <code>description text;</code> |
| Hierarchy Level | <code>[edit interfaces <i>interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</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> |
| 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>Adding an Interface Description to the Configuration</i> • <i>Adding a Logical Unit Description to the Configuration</i> • <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i> • <i>Using DHCP Relay Agent Option 82 Information</i> |

ethernet (Alarm)

| | |
|---------------------------------|---|
| Syntax | <pre>ethernet { link-down (red yellow ignore); }</pre> |
| Hierarchy Level | [edit chassis alarm], [edit chassis interconnect-device <i>name</i> alarm], [edit chassis node-group <i>name</i> alarm] |
| 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 statement is explained separately.— |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Understanding Alarms</i>• <i>Interface Alarm Messages</i> |

ethernet-switching

| | |
|---------------------------------|--|
| Syntax | <pre> ethernet-switching { 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>]; } interface-mode (access trunk); recovery-timeout <i>seconds</i>; storm-control <i>profile-name</i>; vlan { members (<i>vlan-name</i> [<i>-vlan-names</i>] all); } } </pre> |
| Hierarchy Level | [edit interfaces <i>ge-chassis/slot/port unit logical-unit-number</i>] family |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | <p>Configure Ethernet switching protocol family information for the logical interface.</p> <p>The remaining statements are explained separately.</p> |
| Default | You must configure a logical interface to be able to use the physical device. |
| 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 on page 53 • 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);
  configured-flow-control {
    rx-buffers (on | off);
    tx-buffers (on | off);
  }
  (flow-control | no-flow-control);
  link-mode mode;
  (loopback | no-loopback);
  speed (auto-negotiation | no-auto-negotiation);
}
```

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

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

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



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

The statements are explained separately.

Default Enabled.

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

Related Documentation

- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [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 | <ul style="list-style-type: none">• Configuring the Interface Address on page 51 |

family

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

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



```

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

```

```

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

```

| | |
|----------------------------|---|
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [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 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 on page 53](#)
- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- [Configuring Link Aggregation on page 133](#)
- *Configuring IRB Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

fibre-channel (Alarm)


| | |
|---------------------------------|--|
| Syntax | <code>fibre-channel { link-down (red yellow ignore); }</code> |
| Hierarchy Level | [edit chassis alarm], [edit chassis interconnect-device <i>name</i> alarm], [edit chassis node-group <i>name</i> alarm] |
| 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. |
| Related Documentation | <ul style="list-style-type: none">• <i>Understanding Alarms</i>• <i>Interface Alarm Messages</i> |

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 | <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]</p> |
| Release Information | <p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> |
| 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 <i>inet</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet6</i>], and [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>mpls</i>] 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.</p> <p>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.</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>Applying a Filter to an Interface</i> • <i>Junos OS Administration Library for Routing Devices</i> • <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i> • <i>Configuring Firewall Filters (CLI Procedure)</i> |

- *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: 10px 0;"> <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 | <ul style="list-style-type: none"> • configured-flow-control on page 219 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53 |

- *Configuring Gigabit and 10-Gigabit Ethernet Interfaces*
- *Understanding CoS Flow Control (Ethernet PAUSE and PFC)*
- *Junos OS Network Interfaces Library for Routing Devices*



fpc

| | |
|---------------------------------|--|
| Syntax | <pre>fpc slot { auto-speed-detection disable; 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 port-range-high</i> { channel-speed (<i>speed</i> disable-auto-speed-detection); } } }</pre> |
| Hierarchy Level | [edit chassis] |
| Release Information | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | <p>Configure the FPC slot number. For QFX3500 switches, the slot is a line card slot.</p> <p>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.</p> |
| Options | <p>slot—Number of the FPC slot. For QFX3500, QFX3600, QFX5200, and OCX Series devices, the slot number is always 0.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • <i>show chassis fpc</i> • <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i> |

gratuitous-arp-reply

| | |
|---------------------------------|---|
| Syntax | (gratuitous-arp-reply no-gratuitous-arp-reply); |
| Hierarchy Level | [edit interfaces <i>interface-name</i>], [edit interfaces interface-range <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 | <code>hold-time up <i>milliseconds</i> down <i>milliseconds</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-range</i> <i>interface-range-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 10.4R5 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the hold-time value to use to damp shorter interface transitions milliseconds. 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. |
| | <div>  NOTE: <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. </div> |
| | <div>  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. </div> |
| Default | Interface transitions are not damped. |
| Options | down <i>milliseconds</i> —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 Documentation | <ul style="list-style-type: none">• <i>advertise-interval</i>• <i>interfaces (for EX Series switches)</i>• <i>Physical Interface Damping Overview</i>• <i>Damping Shorter Physical Interface Transitions</i>• <i>Damping Longer Physical Interface Transitions</i> |
|------------------------------|--|

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 *interface-name*

| | |
|---------------------------------|--|
| 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. |
| 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 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 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.— |
| 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 on page 53 |

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.— |
| 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 on page 53• Configuring Gigabit and 10-Gigabit Ethernet Interfaces |

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



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

(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the **trunk** option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the **vlan-id** or **vlan-id-list** statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the **access** option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the **vlan-id** statement.



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

| | |
|----------------|---|
| Options | <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> <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</p> |
|----------------|---|

on a single switch. This configuration specifies whether the particular interface assumes the role of interswitch link for the PVLAN domains of which it is a member. This option is supported only on MX240, MX480, and MX960 routers in enhanced LAN mode.

| | |
|---------------------------------|---|
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Configuring Access Mode on a Logical Interface</i>• <i>Configuring a Logical Interface for Trunk Mode</i>• <i>Example: Connecting Access Switches to a Distribution Switch</i>• <i>Tunnel Services Overview</i>• <i>Configuring Tunnel Interfaces on MX Series Routers</i> |

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.

| | |
|---------------------------------|--|
| 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 on page 16 • Interfaces Overview on page 3 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53 • <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

interfaces

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

```

interfaces {
  aex {
    disable;
    aggregated-ether-options {
      configured-flow-control {
        rx-buffers (on | off);
        tx-buffers (on | off);
      }
      (fcoe-lag | no-fcoe-lag);
      (flow-control | no-flow-control);
      lacp mode {
        admin-key key;
        force-up;
        periodic interval;
        system-id mac-address;
      }
      link-speed speed;
      local-bias;
      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);
    }
    (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);
}
(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 | <p>Configure the interfaces on the QFX Series and OCX Series.</p> <p>The following statements and their associated substatements are not supported on OCX Series switches: auto-negotiation, ethernet-switching, fcoe-lag, fibre-channel, fibrechannel-options, mc-ae, speed, vlan, vlan-id, and vlan-tagging</p> <p>Most standard Junos OS configuration statements are available in the Junos OS for a switch. This topic lists Junos OS statements that you commonly use when configuring a switch as well as statements added to support switches only.</p> |
| Options | <p>aex—Configure an aggregated Ethernet interface.</p> <p>xe-0/0/<i>port</i>/—Configure a 10-Gigabit Ethernet interface.</p> <p>ge-0/0/<i>port</i>/—Configure a Gigabit Ethernet interface.</p> <p>fc-0/0/<i>port</i>/—Configure a Fibre Channel interface.</p> <p>meX/—Configure a management interface.</p> <p>mc-ae—Configure a multichassis aggregated Ethernet (MC-AE) interface.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |
| Related Documentation | <ul style="list-style-type: none"> • Interfaces Overview on page 3 • Understanding Interface Ranges on page 16 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces • Configuring Link Aggregation on page 133 • Configuring a Layer 3 Logical Interface on page 118 |


link-down

| | |
|---------------------------------|--|
| Syntax | link-down (red yellow ignore); |
| Hierarchy Level | [edit chassis alarm ethernet], [edit chassis alarm fibre-channel], [edit chassis interconnect-device <i>name</i> alarm ethernet], [edit chassis node-group <i>name</i> alarm fibre-channel] |
| 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 | routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. |

link-mode

| | |
|---------------------------------|--|
| Syntax | <code>link-mode <i>mode</i>;</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. |
| Default | The full-duplex mode is enabled. |
| Options | <p><i>mode</i> —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 on page 53 • <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

link-speed

| | |
|---------------------------------|---|
| Syntax | <code>link-speed <i>speed</i>;</code> |
| Hierarchy Level | [edit interfaces aex aggregated-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 | For aggregated Ethernet interfaces only, set the required link speed. |
| Options | <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 Ethernet bundle. Only link speeds of 40G and 10G are supported. Load balancing will not work if you configure link speeds that are not supported.</p> <p>Aggregated Ethernet links on the QFX Series can have one of the following speed values:</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> NOTE: OCX Series switches only support 10g and 40g interfaces. Mixed rate aggregated Ethernet interfaces are not support on the OCX Series.</p> </div> <ul style="list-style-type: none"> • 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. • 0c192—Links are OC-192. • mixed—Links are 10 Gbps and 40Gbps. |
| 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 Link Aggregation on page 133 |

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

| | |
|---------------------------------|---|
| Syntax | (loopback no-loopback); |
| Hierarchy Level | [edit interfaces <i>interface-name</i> aggregated-ether-options], [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 | For aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode. |
| 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 Ethernet Loopback Capability on page 58 |

mac

| | |
|---------------------------------|--|
| Syntax | mac <i>mac-address</i> ; |
| Hierarchy Level | [edit interfaces <i>interface-name</i>] |
| Release Information | Statement introduced before Junos OS Release 7.4. |
| Description | Set the MAC address of the interface. You can configure the MAC address on the management Ethernet interface (fxp0 or em0) only. |
| Options | mac-address —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0011.2233.4455 or 00:11:22:33:44:55. |
| 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 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 <i>name</i> alarm], [edit chassis node-group <i>name</i> 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 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 Alarms</i>• <i>Interface Alarm Messages</i> |



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 on page 53 • <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i> • Interfaces Overview on page 3 • <i>Interfaces Overview</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

member-range

| | |
|---------------------------------|---|
| Syntax | <code>member-range <i>starting-interface-name ending-interface-name</i>;</code> |
| Hierarchy Level | [edit interfaces <i>interface-range 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 on page 16• Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i>• Interfaces Overview on page 3• <i>Interfaces Overview</i>• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

mtu

| | |
|----------------------------|--|
| Syntax | <code>mtu bytes;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-range</i> <i>interface-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 | <p>Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. On QFX3500, QFX3600, QFX5100, and OCX Series switches, either standalone or as part of the QFabric system, the maximum MTU value on an untagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus four, whereas the maximum MTU value on a tagged packet transiting through an ingress Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight. The maximum MTU value on an untagged or tagged packet transiting through an ingress 10-Gigabit Ethernet interface must be no more than the currently configured MTU value plus eight.</p> <p>Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:</p> <ul style="list-style-type: none"> • 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. <div style="margin-top: 20px;"> <div style="display: flex; align-items: center;">  <div> <p>NOTE: RVIs are not supported on OCX Series switches.</p> </div> </div> <div style="margin-top: 20px;"> <div style="display: flex; align-items: center;">  <div> <p>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.</p> </div> </div> </div> </div> |
| Options | <p>bytes —MTU size.</p> <p>Range: 64 through 9216 bytes</p> |

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 on page 53• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i>• <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</i> |

pic

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

Hierarchy Level [edit chassis fpc *slot*]

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



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

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

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

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

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


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

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

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

Related Documentation • [Channelizing Interfaces on page 63](#)

rx-buffers



| | |
|--|---|
| Syntax | rx-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 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>  <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 | <ul style="list-style-type: none"> • flow-control on page 233 • tx-buffers on page 269 |

- *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 source-address;</code> |
| Hierarchy Level | [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel] |
| Release Information | Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Specify the source address of the tunnel. |
| Default | If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel. |
| Options | <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> |


speed

| | |
|---------------------------------|--|
| Syntax | speed (10g 1g 100m) |
| Hierarchy Level | [edit interfaces <i>interface-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. |
| Description | Use this statement to set the speed of an interface. On 10-Gigabit Ethernet SFP interfaces, autonegotiation is enabled by default and auto-detects the speed to be either 1Gbps or 10Gbps. 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. |
| | <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> |
| | <div>  <p>NOTE: Only 10g and 40g interfaces are supported on OCX Series switches.</p> </div> |
| Default | |
| Options | <ul style="list-style-type: none"> • 10g—10 Gbps • 1g—1 Gbps • 100m—100 Mbps |
| 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"> • auto-negotiation on page 216 • Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53 • <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i> • <i>Junos OS Network Interfaces Library for Routing Devices</i> |


targeted-broadcast

| | |
|---------------------------------|---|
| Syntax | targeted-broadcast; |
| Hierarchy Level | [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] |
| 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 whether the IP packets destined for a Layer 3 broadcast need to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface. |
| Default | When this statement is not included, broadcast packets are sent to the Routing Engine only. |
| 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 IP Directed Broadcast on a Switch</i> • <i>Configuring IP Directed Broadcast (CLI Procedure)</i> • <i>Understanding IP Directed Broadcast</i> • Understanding IP Directed Broadcast on page 107 • Configuring IP Directed Broadcast (CLI Procedure) on page 109 • Example: Configuring IP Directed Broadcast on page 110 |

traceoptions (Individual Interfaces)

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

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 | 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"> • flow-control on page 233 • rx-buffers on page 264 |

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

unit

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

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

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

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

Description



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

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

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

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

Range: 0 through 16,384

The remaining statements are explained separately.

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 on page 53](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces](#)
- [Configuring Link Aggregation on page 133](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

vlan-id

Syntax `vlan-id vlan-id-number;`

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

Release Information 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.



NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

Options *vlan-id-number*—Valid VLAN identifier.
Range: 1 through 4094

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

Related Documentation

- [vlan-tagging on page 273](#)
- [Configuring Gigabit and 10-Gigabit Ethernet Interfaces on page 53](#)
- [Configuring a Layer 3 Logical Interface on page 118](#)
- [Junos OS Network Interfaces Library for Routing Devices](#)

vlan-tagging

| | |
|---------------------------------|---|
| Syntax | vlan-tagging; |
| Hierarchy Level | [edit interfaces <i>interface-name</i>] [edit interfaces interface-range <i>interface-range-name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.3 for the QFX Series. |
| Description | Enable VLAN tagging. The platform receives and forwards single-tag frames with 802.1Q VLAN tags. |
| Default | VLAN tagging is disabled by default. |
| 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-id on page 272• Configuring a Layer 3 Logical Interface on page 118 |

IP Directed Broadcast Configuration Statement

- [targeted-broadcast](#) on page 275

targeted-broadcast

| | |
|---------------------------------|---|
| Syntax | targeted-broadcast; |
| Hierarchy Level | [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] |
| 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 whether the IP packets destined for a Layer 3 broadcast need to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface. |
| Default | When this statement is not included, broadcast packets are sent to the Routing Engine only. |
| 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 IP Directed Broadcast on a Switch</i> • <i>Configuring IP Directed Broadcast (CLI Procedure)</i> • <i>Understanding IP Directed Broadcast</i> • Understanding IP Directed Broadcast on page 107 • Configuring IP Directed Broadcast (CLI Procedure) on page 109 • Example: Configuring IP Directed Broadcast on page 110 |

CHAPTER 16

LAGs and LACP Configuration Statements

- [aggregated-devices on page 278](#)
- [aggregated-ether-options on page 279](#)
- [chassis on page 281](#)
- [802.3ad on page 282](#)
- [device-count on page 283](#)
- [ethernet on page 283](#)
- [force-up on page 284](#)
- [lacp \(802.3ad\) on page 285](#)
- [lacp \(Aggregated Ethernet\) on page 286](#)
- [periodic on page 287](#)

aggregated-devices

| | |
|---------------------------------|--|
| Syntax | <pre>aggregated-devices { ethernet { device-count <i>number</i>; } }</pre> |
| 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. Statement introduced in Junos OS Release 14.2R3 |
| Description | Configure properties for aggregated devices on the switch. The remaining statements are explained separately. |
| Default | Aggregated devices are disabled. |
| Required Privilege Level | interface—To view this statement in the configuration. interface-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• <i>Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion</i>• Understanding Aggregated Ethernet Interfaces and LACP on page 129• Configuring Link Aggregation on page 133• <i>Configuring Link Aggregation</i>• Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136• <i>Junos OS Network Interfaces Library for Routing Devices</i> |

aggregated-ether-options

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

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

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

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

Description Configure properties specific to a specific aggregated Ethernet interface.



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



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



NOTE: The **resilient-hash** statement is not supported on QFX10002 switches.

The statements are explained separately.

Default Options are not enabled.

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

Related Documentation

- [Understanding Aggregated Ethernet Interfaces and LACP on page 129](#)
- [Configuring Aggregated Ethernet LACP on page 132](#)
- [Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 140](#)
- *Junos OS Network Interfaces Library for Routing Devices*

chassis

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

Hierarchy Level [edit]

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

Description Configure chassis-specific properties for the switch.

The remaining statements are explained separately.

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

- Related Documentation**
- [Understanding Link Aggregation and Link Aggregation Control Protocol in a Junos Fusion](#)
 - [Configuring Link Aggregation](#)
 - [Configuring Link Aggregation on page 133](#)

802.3ad

Syntax

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

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

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

Description Specify the aggregated Ethernet logical interface number.



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



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

Options aex—Aggregated Ethernet logical interface number.

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

- Related Documentation**
- [Configuring Link Aggregation on page 133](#)
 - [Configuring Aggregated Ethernet LACP on page 132](#)
 - [Understanding Aggregated Ethernet Interfaces and LACP on page 129](#)
 - [Troubleshooting an Aggregated Ethernet Interface on page 146](#)
 - [Junos OS Network Interfaces Library for Routing Devices](#)


device-count

| | |
|---------------------------------|--|
| Syntax | <code>device-count <i>number</i>;</code> |
| Hierarchy Level | [edit chassis aggregated-devices ethernet], [edit chassis node-group <i>name</i> aggregated-devices ethernet] |
| 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 the number of aggregated Ethernet logical devices available to the switch. |
| 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> • <i>Configuring Link Aggregation</i> • Configuring Link Aggregation on page 133 • Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136 |



ethernet

| | |
|---------------------------------|--|
| Syntax | <code>ethernet { device-count <i>number</i>; }</code> |
| 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. |
| 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> • <i>Configuring Link Aggregation</i> • Configuring Link Aggregation on page 133 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

force-up

| | |
|--|--|
| Syntax | force-up; |
| Hierarchy Level | [edit interfaces interface-name ether-options 802.3ad lacp ; [edit interfaces interface-name aggregated-ether-options lacp ; |
| 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 state of the interface as up when the peer has limited LACP capability. You can also configure the peer interface (in MC-LAG) to remain up even with limited LACP capability. |
| <div>  NOTE: The force-up option is not supported on QFX10002 switches. </div> | |
| 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 Link Aggregation • Understanding Aggregated Ethernet Interfaces and LACP on page 129 • Configuring Aggregated Ethernet LACP on page 132 • Example: Configuring Link Aggregation with LACP Between a QFX Series Product and an Aggregation Switch on page 140 • Junos OS Network Interfaces Library for Routing Devices • |

lcp (802.3ad)

| | |
|---|---|
| Syntax | <pre>lcp { force-up; (primary backup); port-priority; }</pre> |
| Hierarchy Level | [edit interfaces interface-name ether-options 802.3ad] |
| Release Information | Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series. |
| Description | Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately. |
| <div>  <p>NOTE: The port-priority statement is not supported on QFabric systems.</p> </div> | |
| <div>  <p>NOTE: The force-up statement is not supported on QFX10002 switches.</p> </div> | |
| 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 Link Aggregation • Configuring Link Aggregation on page 133 • Configuring Aggregated Ethernet LACP on page 132 • Understanding Aggregated Ethernet Interfaces and LACP on page 129 |

lACP (Aggregated Ethernet)

Syntax 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;
 }

Hierarchy Level [edit [interfaces interface-name aggregated-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 Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces. The remaining statement is explained separately.



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

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](#)
- [Configuring Link Aggregation on page 133](#)
- [Configuring Aggregated Ethernet LACP on page 132](#)
- [Configuring LACP Link Protection of Aggregated Ethernet Interfaces \(CLI Procedure\)](#)
- [Understanding Aggregated Ethernet Interfaces and LACP on page 129](#)

periodic

| | |
|---------------------------------|--|
| Syntax | <code>periodic (fast slow);</code> |
| Hierarchy Level | [edit interfaces <code>aex</code> aggregated-ether-options <code>lACP</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 | Configure the interval for periodic transmission of LACP packets. |
| Default | <code>fast</code> |
| Options | <p><i>interval</i>—Interval at which to periodically transmit LACP packets:</p> <ul style="list-style-type: none"> • <i>fast</i>—Receive packets every second. This is the default. • <i>slow</i>—Receive packets every 30 seconds. |
| 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 Link Aggregation</i> • Configuring Aggregated Ethernet LACP on page 132 • Understanding Aggregated Ethernet Interfaces and LACP on page 129 • <i>Junos OS Network Interfaces Library for Routing Devices</i> |

Local Link Bias Configuration Statements

- [local-bias on page 289](#)

local-bias

| | |
|---------------------------------|--|
| Syntax | local-bias; |
| Hierarchy Level | [edit interfaces aex aggregated-ether-options] |
| Release Information | Statement introduced in Junos OS Release 13.2X51-D20 for EX Series switches and QFX Series devices. |
| Description | <p>Enable local link bias for all links in the aggregated Ethernet interface.</p> <p>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.</p> <p>You should enable local link bias if you want to conserve VCP bandwidth by always forwarding egress unicast traffic on a LAG bundle out of a local link. You should not enable local link bias if you want egress traffic load-balanced as it exits the Virtual Chassis or VCF.</p> |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Local Link Bias (CLI Procedure) on page 125 • Understanding Local Link Bias on page 123 |

CHAPTER 18

Redundant Trunk Groups Configuration Statements

- [group \(Redundant Trunk Groups\) on page 292](#)
- [interface \(Redundant Trunk Groups\) on page 293](#)
- [preempt-cutover-timer on page 294](#)
- [redundant-trunk-group on page 295](#)

group (Redundant Trunk Groups)

| | |
|---------------------------------|---|
| Syntax | <pre>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: [edit switch-options redundant-trunk-group]For platforms without ELS: [edit ethernet-switching-options redundant-trunk-group] |
| 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>Getting Started with Enhanced Layer 2 Software</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: The group name must be a string “rtgn” where n is a number from 0 through 15, such as “rtg2” or “rtg10”.For platforms without ELS: The group name must start with a letter and can consist of letters, numbers, dashes, and underscores. <p>The remaining statements are explained separately.</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</i>Example: Configuring Redundant Trunk Links for Faster Recovery on page 153Understanding Redundant Trunk Links on page 151 |

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>Getting Started with Enhanced Layer 2 Software</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</i> • Example: Configuring Redundant Trunk Links for Faster Recovery on page 153 • Understanding Redundant Trunk Links on page 151 |

preempt-cutover-timer

| | |
|---------------------------------|--|
| Syntax | <code>preempt-cutover-timer seconds;</code> |
| Hierarchy Level | <ul style="list-style-type: none">For platforms with ELS: [edit switch-options redundant-trunk-group <i>group name</i>]For platforms without ELS: [edit ethernet-switching-options redundant-trunk-group <i>group name</i>] |
| Release Information | Statement introduced in Junos OS Release 11.1 for EX Series switches. Hierarchy level [edit switch-options] introduced in Junos OS Release 13.2X50-D10 (ELS). (See <i>Getting Started with Enhanced Layer 2 Software</i> for information about ELS.) Statement introduced in Junos OS Release 13.2X50-D15 for the QFX Series. |
| 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. |
| 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 | seconds —Number of seconds that the primary link waits to take over from the active secondary link. Range: 1 through 600 seconds |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"><i>Example: Configuring Redundant Trunk Links for Faster Recovery</i>Example: Configuring Redundant Trunk Links for Faster Recovery on page 153Understanding Redundant Trunk Links on page 151 |

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: [edit switch-options] For platforms without ELS: [edit ethernet-switching-options] |
| 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>Getting Started with Enhanced Layer 2 Software</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.</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</i> Example: Configuring Redundant Trunk Links for Faster Recovery on page 153 Understanding Redundant Trunk Links on page 151 |

CHAPTER 19

Resilient Hashing Configuration Statements

- [ecmp-resilient-hash](#) on page 297
- [enhanced-hash-key](#) on page 298
- [hash-mode](#) on page 301
- [inet \(enhanced-hash-key\)](#) on page 303
- [inet6 \(enhanced-hash-key\)](#) on page 305
- [resilient-hash](#) on page 307

[ecmp-resilient-hash](#)

| | |
|----------------------------|---|
| Syntax | <code>ecmp-resilient-hash;</code> |
| Hierarchy Level | [edit forwarding-options enhanced-hash-key] |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series. |
| Description | Enable resilient hashing for ECMP groups, to minimize remapping of destination paths. |



NOTE:

| | |
|---------------------------------|---|
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Configuring Resilient Hashing for Trunk/ECMP Groups on page 171 |

enhanced-hash-key

List of Syntax [Syntax \(EX Series and QFX5100 Switch\) on page 298](#)
[Syntax \(QFX10002 and QFX10008 Switches\) on page 298](#)

Syntax (EX Series and QFX5100 Switch)

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

Syntax (QFX10002 and QFX10008 Switches)

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



```

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

```

| | |
|---------------------------------|--|
| Hierarchy Level | [edit forwarding-options] |
| Release Information | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p>The fabric-load-balance statement introduced in Junos OS Release 14.1X53-D10.</p> <p>The hash-seed statement introduced in Junos OS Release 15.1X53-D30.</p> |
| Description | <p>Configure the hashing key used to hash link aggregation group (LAG) and equal-cost multipath (ECMP) traffic, or enable adaptive load balancing (ALB) in a Virtual Chassis Fabric (VCF).</p> <p>The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.</p> <p>For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.</p> <p>On QFX10002 and QFX 10008 switches, you can configure the hash seed for load balancing.</p> <p>By default, the QFX10002 and QFX10008 switches use the system MAC address to generate a hash seed value. You can configure the hash seed value using the hash-seed statement at the [edit forwarding-options enhanced-hash-key] hierarchy level. Set a value between 0 and 4294967295. If you do not configure a hash seed value, the system will by generate a hash seed value based on the system MAC address.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p> |

**Related
Documentation**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 169](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 161](#)

hash-mode

| | |
|---------------------------------|--|
| Syntax | <pre>hash-mode { layer2-header; layer2-payload; }</pre> |
| Hierarchy Level | [edit forwarding-options enhanced-hash-key] |
| Release Information | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p>Statement is not supported on QFX10002 and QFX 10008 switches.</p> |
| Description | <p>Select the mode for the hashing algorithm.</p> <p>The hashing algorithm is used to make traffic-forwarding decisions for traffic entering a LAG bundle or for traffic exiting a switch when ECMP is enabled.</p> <p>For LAG bundles, the hashing algorithm determines how traffic entering a LAG bundle is placed onto the bundle's member links. The hashing algorithm tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>When ECMP is enabled, the hashing algorithm determines how incoming traffic is forwarded to the next-hop device.</p> <p>The hash mode that is set using this statement determines which fields are inspected by the hashing algorithm. You must set the hash mode to layer2-payload if you want the hashing algorithm to inspect fields in the Layer 2 payload when making hashing decisions. You must set the hash mode to layer2-header if you want the hashing algorithm to inspect fields in the Layer 2 header when making hashing decisions.</p> <p>If the hash mode is set to layer2-payload, you can set the fields used by the hashing algorithm to hash IPv4 traffic using the set forwarding-options enhanced-hash-key inet statement. You can set the fields used by the hashing algorithm to hash IPv6 traffic using the set forwarding-options enhanced-hash-key inet6 statement.</p> <p>If the hash mode is set to layer2-header, you can set the fields that the hashing algorithm inspects in the Layer 2 header using the set forwarding-options enhanced-hash-key layer2 statement.</p> |
| Default | layer2-payload |
| Options | <p>layer-2-payload—Set the hashing algorithm to use fields in the Layer 2 payload to make hashing decisions.</p> <p>layer-2-header—Set the hashing algorithm to use fields in the Layer 2 header to make hashing decisions.</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**

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 169](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 161](#)
- [enhanced-hash-key on page 298](#)
- [inet on page 303](#)
- [inet6 on page 305](#)
- [layer2](#)

inet (enhanced-hash-key)

| | |
|--|--|
| Syntax (EX Series and QFX5100 Switch) | <pre>inet { no-ipv4-destination-address; no-ipv4-source-address; no-l4-destination-port; no-l4-source-port; no-protocol; vlan-id; }</pre> |
| Syntax (QFX10002 and QFX10008 Switches) | <pre>inet { no-ipv4-destination-address; no-ipv4-source-address; no-l4-destination-port; no-l4-source-port; }</pre> |
| Hierarchy Level | [edit forwarding-options enhanced-hash-key] |
| Release Information | <p>Statement introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and QFX10008 Switches.</p> |
| Description | <p>Select the payload fields in IPv4 traffic used by the hashing algorithm to make hashing decisions.</p> <p>When IPv4 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using the inet statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or how to forward traffic to the next hop device when ECMP is enabled.</p> <p>The hashing algorithm, when used to hash LAG bundle traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>The hashing algorithm only inspects the IPv4 fields in the payload to make hashing decisions when the hash mode is set to layer2-payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the set forwarding-options enhanced-hash-key hash-mode layer2-payload statement.</p> |
| Default | <p>The following fields are used by the hashing algorithm to make hashing decisions for IPv4 traffic:</p> <ul style="list-style-type: none"> • IP destination address • IP source address • Layer 4 destination port • Layer 4 source port |

- Protocol

| | |
|----------------|--|
| Options | no-ipv4-destination-address —Exclude the IPv4 destination address field from the hashing algorithm. |
| | no-ipv4-source-address —Exclude the IPv4 source address field from the hashing algorithm. |
| | no-l4-destination-port —Exclude the Layer 4 destination port field from the hashing algorithm. |
| | no-l4-source-port —Exclude the Layer 4 source port field from the hashing algorithm. |
| | no-protocol —Exclude the protocol field from the hashing algorithm. |
| | vlan-id —Include the VLAN ID field in the hashing algorithm. |

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

| | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none">• Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 169• Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 161• <i>Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)</i>• <i>hash-seed</i>• enhanced-hash-key on page 298• hash-mode on page 301• inet6 on page 305 |
|------------------------------|---|

inet6 (enhanced-hash-key)

| | |
|--|--|
| List of Syntax | Syntax (EX Series and QFX5100 Switch) on page 305 Syntax (QFX10002 and QFX10008 Switches) on page 305 |
| Syntax (EX Series and QFX5100 Switch) | <pre>inet6 { no-ipv6-destination-address; no-ipv6-source-address; no-l4-destination-port; no-l4-source-port; no-next-header; vlan-id; }</pre> |
| Syntax (QFX10002 and QFX10008 Switches) | <pre>inet6 { ipv6-flow-label; no-ipv6-destination-address; no-ipv6-source-address; no-l4-destination-port; no-l4-source-port; }</pre> |
| Hierarchy Level | [edit forwarding-options enhanced-hash-key] |
| Release Information | <p>Statement introduced in Junos OS Release 13.2X51-D15 on EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2X51-D20 on QFX Series devices.</p> <p>Statement introduced in Junos OS Release 15.1X53-D30 on QFX10002 and QFX 10008 switches.</p> |
| Description | <p>Select the payload fields in an IPv6 packet used by the hashing algorithm to make hashing decisions.</p> <p>When IPv6 traffic enters a LAG and the hash mode is set to Layer 2 payload, the hashing algorithm checks the fields configured using this statement and uses the information in the fields to decide how to place traffic onto the LAG bundle's member links or to forward traffic to the next hop device when ECMP is enabled.</p> <p>The hashing algorithm, when used to hash LAG traffic, always tries to manage bandwidth by evenly load-balancing all incoming traffic across the member links in the bundle.</p> <p>The hashing algorithm only inspects the IPv6 fields in the payload to make hashing decisions when the hash mode is set to Layer 2 payload. The hash mode is set to Layer 2 payload by default. You can set the hash mode to Layer 2 payload using the set forwarding-options enhanced-hash-key hash-mode layer2-payload statement.</p> |
| Default | <p>The data in the following fields are used by the hashing algorithm to make hashing decisions for IPv6 traffic:</p> <ul style="list-style-type: none"> • IP destination address • IP source address • Layer 4 destination port |

- Layer 4 source port
- Next header

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

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

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

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

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

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

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

Related Documentation

- [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\) on page 169](#)
- [Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 161](#)
- *Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic (QFX 10002 and QFX 10008 Switches)*
- *hash-seed*
- [enhanced-hash-key on page 298](#)
- [hash-mode on page 301](#)
- [inet on page 303](#)

resilient-hash

| | |
|----------------------------|--|
| Syntax | resilient-hash; |
| Hierarchy Level | [edit interfaces aex aggregated-ether-options]] |
| Release Information | Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series. |



NOTE: Configuring resilient hashing on LAGs is not supported on QFX10002 and QFX 10008 switches.

| | |
|---------------------------------|---|
| Description | Enable resilient hashing for a LAG, to minimize remapping of destination paths. |
| Required Privilege Level | system—To view this statement in the configuration. system-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none"> • Configuring Resilient Hashing for Trunk/ECMP Groups on page 171 |

CHAPTER 20

Uplink Failure Detection Configuration Statements

- [group](#) on page 309
- [link-to-disable](#) on page 310
- [link-to-monitor](#) on page 310
- [uplink-failure-detection](#) on page 311

group

Syntax `group group-name {
 link-to-monitor interface-name;
 link-to-disable interface-name;
 }`

Hierarchy Level `[edit protocols uplink-failure-detection]`

Release Information Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Configure a group of uplink and downlink interfaces for uplink failure detection.

Options ***group-name***—Name of the uplink failure detection group.

The remaining statements are explained separately.

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

Related Documentation

- [Overview of Uplink Failure Detection](#) on page 175
- [Configuring Interfaces for Uplink Failure Detection](#) on page 177
- [Example: Configuring Interfaces for Uplink Failure Detection](#) on page 178

link-to-disable

| | |
|---------------------------------|---|
| Syntax | <code>link-to-disable <i>interface-name</i>;</code> |
| Hierarchy Level | <code>[edit protocols uplink-failure-detection group <i>group-name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of eight downlink interfaces in a group. |
| Options | <i>interface-name</i> —Name of the downlink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Overview of Uplink Failure Detection on page 175• Configuring Interfaces for Uplink Failure Detection on page 177• Example: Configuring Interfaces for Uplink Failure Detection on page 178 |

link-to-monitor

| | |
|---------------------------------|---|
| Syntax | <code>link-to-monitor <i>interface-name</i>;</code> |
| Hierarchy Level | <code>[edit protocols uplink-failure-detection group <i>group-name</i>]</code> |
| Release Information | Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of eight uplink interfaces in a group. |
| Options | <i>interface-name</i> —Name of the uplink interface in an uplink failure detection group. The interface can be a physical interface or a logical interface. |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Overview of Uplink Failure Detection on page 175• Configuring Interfaces for Uplink Failure Detection on page 177• Example: Configuring Interfaces for Uplink Failure Detection on page 178 |

uplink-failure-detection

| | |
|---------------------------------|---|
| Syntax | <pre>uplink-failure-detection { group <i>group-name</i> { link-to-monitor <i>interface-name</i>; link-to-disable <i>interface-name</i>; } }</pre> |
| Hierarchy Level | [edit protocols] |
| Release Information | Statement introduced in Junos OS Release 12.1 for the QFX Series. |
| Description | <p>Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failure information to the downlink interfaces.</p> <p>The remaining statements are explained separately.</p> |
| Required Privilege Level | admin—To view this statement in the configuration. admin-control—To add this statement to the configuration. |
| Related Documentation | <ul style="list-style-type: none">• Overview of Uplink Failure Detection on page 175• Configuring Interfaces for Uplink Failure Detection on page 177• Example: Configuring Interfaces for Uplink Failure Detection on page 178 |

CHAPTER 21

Interfaces Operational Commands

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

monitor interface

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

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

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



NOTE: This command is not supported on the QFX3000 QFabric 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 32 on page 314](#). The keys are not case-sensitive.

Table 32: 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. |

Table 32: Output Control Keys for the monitor interface interface-name Command (*continued*)

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

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in [Table 33 on page 315](#). The keys are not case-sensitive.

Table 33: 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 317](#)
- [monitor interface \(OTN Interface\) on page 318](#)
- [monitor interface \(MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface\) on page 319](#)
- [monitor interface \(MX480 Router with MPC5E and 100-Gigabit Ethernet Interface\) on page 320](#)
- [monitor interface \(MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 320](#)
- [monitor interface \(MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface\) on page 321](#)
- [monitor interface \(MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface\) on page 322](#)
- [monitor interface \(Logical\) on page 322](#)
- [monitor interface \(QFX3500 Switch\) on page 323](#)

[monitor interface traffic on page 323](#)

[monitor interface traffic \(QFX3500 Switch\) on page 324](#)

[monitor interface traffic detail \(QFX3500 Switch\) on page 324](#)

Output Fields [Table 34 on page 316](#) describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 34: 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 |
| Remote Statistics | (Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. <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 34: monitor interface Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------|--|-----------------|
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize.</p> <ul style="list-style-type: none"> • 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)
    Input bytes:                  6290065 (0 bps)
    Output packets:               10376 (0 pps)
    Output bytes:                 10365540 (0 bps)
Encapsulation statistics:
    Input keepalives:             1901
    Output keepalives:           1901
    NCP state: Opened
    LCP state: Opened
Error statistics:
    Input errors:                 0
    Input drops:                 0
    Input framing errors:        0
    Policed discards:            0
    L3 incompletes:              0
    L2 channel errors:           0
    L2 mismatch timeouts:        0
    Carrier transitions:          1
    Output errors:               0
    Output drops:               0
    Aged packets:                0
Active alarms : None
Active defects: None
SONET error counts/seconds:
    LOS count                    1
    LOF count                    1
    SEF count                    1
    ES-S                        0
    SES-S                        0
SONET statistics:
    BIP-B1                      458871

```

```

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

```

Received SONET overhead:

```

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

```

Transmitted SONET overhead:

```

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

```

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

monitor interface (OTN Interface)

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

```

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

```

```

LOF 9
OTN OTU - FEC Statistics
  Corr err ratio N/A
  Corr bytes 0
  Uncorr words 0
OTN OTU - Counters
  BIP 0
  BBE 0
  ES 0
  SES 0
  UAS 422
OTN ODU - Counters
  BIP 0
  BBE 0
  ES 0
  SES 0
  UAS 422
OTN ODU - Received Overhead 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)
  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
PCS statistics:
  Bit Errors 0
  Errored blocks 4
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

```

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 [0]
BBE 0 [0]
ES 24 [0]
SES 0 [0]
UAS 1256 [0]
OTN ODU - Received Overhead
APSPCC 0-3: 00 00 00 00 [0]

```

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:
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)                [0]
Output bytes:                0 (0 bps)                [0]
Input packets:              0 (0 pps)                [0]
Output packets:             0 (0 pps)                [0]
Error statistics:
Input errors:                0                        [0]
Input drops:                 0                        [0]
Input framing errors:        0                        [0]
Policed discards:            0                        [0]
L3 incompletes:              0                        [0]
L2 channel errors:           0                        [0]
L2 mismatch timeouts:        0                        [0]
Carrier transitions:          1                        [0]
Output errors:               0                        [0]
Output drops:                0                        [0]
Aged packets:                0                        [0]

```

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

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

```

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

Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:                               Current delta
  Input bytes:                                0 (0 bps)                                [0]
  Output bytes:                               0 (0 bps)                                [0]
  Input packets:                              0 (0 pps)                                [0]
  Output packets:                             0 (0 pps)                                [0]
Error statistics:
  Input errors:                               0                                [0]
  Input drops:                                0                                [0]
  Input framing errors:                       0                                [0]
  Policed discards:                           0                                [0]
  L3 incompletes:                             0                                [0]
  L2 channel errors:                          0                                [0]
  L2 mismatch timeouts:                       0                                [0]
  Carrier transitions:                         3                                [0]
  Output errors:                              0                                [0]
  Output drops:                               0                                [0]
  Aged packets:                               0                                [0]
OTN Link 0
OTN Alarms:
OTN Defects:
OTN OC - Seconds
  LOS                                         0                                [0]
  LOF                                         0                                [0]
OTN OTU - FEC Statistics
  Corr err ratio                             N/A
  Corr bytes                                 0                                [0]
  Uncorr words                               0                                [0]
OTN OTU - Counters
  BIP                                         0                                [0]
  BBE                                         0                                [0]
  ES                                          0                                [0]
  SES                                         0                                [0]
  UAS                                         0                                [0]
OTN ODU - Counters
  BIP                                         0                                [0]
  BBE                                         0                                [0]
  ES                                          0                                [0]
  SES                                         0                                [0]
  UAS                                         0                                [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                      [0]
  Output bytes:                 0                      [0]
  Input packets:                0                      [0]
  Output packets:               0                      [0]
Remote statistics:
  Input bytes:                  0 (0 bps)              [0]
  Output bytes:                 0 (0 bps)              [0]
  Input packets:                0 (0 pps)              [0]
  Output packets:               0 (0 pps)              [0]
Traffic statistics:
  Destination address: 192.168.8.193, Local: 192.168.8.21

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

```

monitor interface (QFX3500 Switch)

```

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

Interface warnings:
  o Outstanding LINK alarm

```

monitor interface traffic

```

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

Interface  Link  Input packets  (pps)  Output packets  (pps)
so-1/0/0   Down    0              (0)     0              (0)
so-1/1/0   Down    0              (0)     0              (0)
so-1/1/1   Down    0              (0)     0              (0)
so-1/1/2   Down    0              (0)     0              (0)

```

| | | | | | |
|----------|------|--------|-----|-------|-----|
| so-1/1/3 | Down | 0 | (0) | 0 | (0) |
| t3-1/2/0 | Down | 0 | (0) | 0 | (0) |
| t3-1/2/1 | Down | 0 | (0) | 0 | (0) |
| t3-1/2/2 | Down | 0 | (0) | 0 | (0) |
| t3-1/2/3 | Down | 0 | (0) | 0 | (0) |
| so-2/0/0 | Up | 211035 | (1) | 36778 | (0) |
| so-2/0/1 | Up | 192753 | (1) | 36782 | (0) |
| so-2/0/2 | Up | 211020 | (1) | 36779 | (0) |
| so-2/0/3 | Up | 211029 | (1) | 36776 | (0) |
| so-2/1/0 | Up | 189378 | (1) | 36349 | (0) |
| so-2/1/1 | Down | 0 | (0) | 18747 | (0) |
| so-2/1/2 | Down | 0 | (0) | 16078 | (0) |
| so-2/1/3 | Up | 0 | (0) | 80338 | (0) |
| at-2/3/0 | Up | 0 | (0) | 0 | (0) |
| at-2/3/1 | Down | 0 | (0) | 0 | (0) |

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

monitor interface traffic (QFX3500 Switch)

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

| Interface | Link | Input packets | (pps) | Output packets | (pps) |
|-----------|------|---------------|-------|----------------|-------|
| ge-0/0/0 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/1 | Up | 392187 | (0) | 392170 | (0) |
| ge-0/0/2 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/3 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/4 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/5 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/6 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/7 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/8 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/9 | Up | 392184 | (0) | 392171 | (0) |
| ge-0/0/10 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/11 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/12 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/13 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/14 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/15 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/16 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/17 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/18 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/19 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/20 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/21 | Down | 0 | (0) | 0 | (0) |
| ge-0/0/22 | Up | 392172 | (0) | 392187 | (0) |
| ge-0/0/23 | Up | 392185 | (0) | 392173 | (0) |
| vcp-0 | Down | 0 | | 0 | |
| vcp-1 | Down | 0 | | 0 | |
| ae0 | Down | 0 | (0) | 0 | (0) |
| bme0 | Up | 0 | | 1568706 | |

monitor interface traffic detail (QFX3500 Switch)

```
user@switch> monitor interface traffic detail
switch                                     Seconds: 74                               Time: 16:03:02
```

| Interface Description | Link | Input packets | (pps) | Output packets | (pps) |
|--------------------------|------|---------------|-------|----------------|-------|
|--------------------------|------|---------------|-------|----------------|-------|

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

show interfaces diagnostics optics

| | |
|---------------------------------|--|
| 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. For a list of transceivers supported on EX Series switches and their specifications, including DOM support, see <i>Pluggable Transceivers Supported on EX Series Switches</i>.</p> <p>Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.</p> |
| 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 79 • <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 333</p> <p>show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 334</p> <p>show interfaces diagnostics optics xe-0/1/0 (XFP Transceiver) on page 335</p> <p>show interfaces diagnostics optics et-3/0/0 (QSFP+ Transceiver) on page 336</p> <p>show interfaces diagnostics optics et-4/1/0 (CFP Transceiver) on page 337</p> |
| Output Fields | Table 35 on page 326 lists the output fields for the show interfaces diagnostics optics command. Output fields are listed in the approximate order in which they appear. |

Table 35: show interfaces diagnostics optics Output Fields

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

Table 35: show interfaces diagnostics optics Output Fields (*continued*)

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

Table 35: show interfaces diagnostics optics Output Fields (*continued*)

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

Table 35: show interfaces diagnostics optics Output Fields (*continued*)

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

Table 35: show interfaces diagnostics optics Output Fields (*continued*)

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

Table 35: show interfaces diagnostics optics Output Fields (*continued*)

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

Table 35: show interfaces diagnostics optics Output Fields (*continued*)

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

Table 35: show interfaces diagnostics optics Output Fields (*continued*)

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

Sample Output

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

```

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

```

```

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

```

Sample Output

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

```

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

```

```

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

```

Sample Output

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

```

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

```

Sample Output

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

```

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

```

Sample Output

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

```

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

```

```

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

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

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

```



```

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

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

```

show interfaces ge

| | |
|---------------------------------|---|
| Syntax | <code>show interfaces <i>device-name:type-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><routing-instance (all <i>instance-name</i>)></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code> |
| Release Information | 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. |
| Options | <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p><i>device-name:type-fpc/pic/port</i>—The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name can contain a maximum of 128 characters and cannot contain any colons.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>routing instance (all <i>instance-name</i>)—(Optional) Display the name of an individual routing-instance or display all routing-instances.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none">• Monitoring Interface Status and Traffic on page 79• Troubleshooting Network Interfaces on page 80• Troubleshooting an Aggregated Ethernet Interface on page 146• Junos OS Network Interfaces Library for Routing Devices |
| List of Sample Output | show interfaces on page 348 show interfaces brief on page 348 show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled) on page 348 show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled) on page 349 |

[show interfaces extensive \(Symmetric Flow Control and Autonegotiation Enabled\) on page 350](#)

[show interfaces extensive \(Asymmetric Flow Control and Autonegotiation Enabled\) on page 352](#)

[show interfaces terse on page 354](#)

[show interfaces terse \(QFabric Systems\) on page 354](#)

Output Fields Table 36 on page 341 lists the output fields for the **show interfaces ge** command. Output fields are listed in the approximate order in which they appear.

Table 36: show interfaces ge Output Fields

| Field Name | Field Description | Level of Output |
|--------------------------------|--|-------------------------------|
| Physical Interface | | |
| Physical interface | Name of the physical interface. | All levels |
| Enabled | State of the interface: Enabled or Disabled . | All levels |
| Interface index | Index number of the physical interface, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP index number for the physical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Description | Optional user-specified description. | brief detail extensive |
| Link-level type | Encapsulation being used on the physical interface. | All levels |
| MTU | Maximum transmission unit size on the physical interface. The default is 1514. | All levels |
| Speed | Speed at which the interface is running. | All levels |
| Loopback | Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote . | All levels |
| Source filtering | Source filtering status: Enabled or Disabled . | All levels |
| Flow control | Flow control status: Enabled or Disabled . <i>NOTE:</i> This field is only displayed if asymmetric flow control is not configured. | detail extensive |
| 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. <i>NOTE:</i> This field is only displayed if asymmetric flow control is configured. | detail extensive |

Table 36: show interfaces ge Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| 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 |
| CoS queues | Number of CoS queues configured. | detail extensive none |
| Hold-times | Current interface hold-time up and hold-time down, in milliseconds. | detail extensive |
| Current address | Configured MAC address. | detail extensive none |
| Hardware address | MAC address of the hardware. | detail extensive none |
| Last flapped | Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago) . For example, Last flapped: 2008-01-16 10:52:40 UTC (3d 22:58 ago) . | detail extensive none |
| Statistics last cleared | Time when the statistics for the interface were last set to zero. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the physical interface. <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 |

Table 36: show interfaces ge Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------|---|------------------|
| 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 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 36: 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 |
| 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 |
| 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 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 36: 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 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 runts (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 36: 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"> • 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 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 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: <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 PAUSE on both receive and transmit or PAUSE only on receive). For asymmetric PAUSE, shows if the PAUSE transmit and PAUSE receive states on the interface are 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 |
| 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. | extensive |
| Logical Interface | | |
| Logical interface | Name of the logical interface. | All levels |
| Index | Index number of the logical interface, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | SNMP interface index number for the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Flags | Information about the logical interface. | All levels |

Table 36: show interfaces ge Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|---|------------------------------|
| 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 statistics 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 |
| 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 |
| Destination | IP address of the remote side of the connection. | detail extensive none |
| Local | IP address of the logical interface. | detail extensive none |
| Broadcast | Broadcast address of the logical interlace. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces

```

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

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

```

show interfaces brief

```

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

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

```

show interfaces detail (Symmetric Flow Control and Autonegotiation Enabled)

```

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

```

```

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

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

Active alarms : None
Active defects : None

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

```

show interfaces detail (Asymmetric Flow Control and Autonegotiation Enabled)

```

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

```

```

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

```

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

```

Active alarms : None
Active defects : None

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

```

show interfaces extensive (Symmetric Flow Control and Autonegotiation Enabled)

```

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

```

```

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

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

```

| | Queued packets | Transmitted packets | Dropped packets |
|----------------|----------------|---------------------|-----------------|
| 0 best-effort | 0 | 0 | 0 |
| 2 no-loss | 0 | 0 | 0 |
| 3 fcoe | 0 | 0 | 0 |
| 7 network-cont | 0 | 0 | 0 |

```

Queue number:      Mapped forwarding classes
0                  best-effort
2                  no-loss
3                  fcoe
7                  network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:

```

| | Receive | Transmit |
|--------------------|---------|----------|
| Total octets | 0 | 0 |
| Total packets | 0 | 0 |
| Unicast packets | 0 | 0 |
| Broadcast packets | 0 | 0 |
| Multicast packets | 0 | 0 |
| CRC/Align errors | 0 | 0 |
| FIFO errors | 0 | 0 |
| MAC control frames | 0 | 0 |
| MAC pause frames | 0 | 0 |
| Oversized frames | 0 | |
| Jabber frames | 0 | |
| Fragment frames | 0 | |

```

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

```

show interfaces extensive (Asymmetric Flow Control and Autonegotiation Enabled)

```

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

```

```

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

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

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

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

```

```

Output packet count                                0
Output packet pad count                            0
Output packet error count                          0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link Partner:
  Link mode: Full-duplex, Flow control: None, Remote fault: OK,
  Link partner Speed: 1000 Mbps
Local resolution:
  Flow control: enable PAUSE transmit and Disable PAUSE receive, Remote
fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
0 best-effort              75      750000000  75      0      low
none
7 network-control          5       50000000  5       0      low
none
8 mcast-be                 15     150000000  15      0      low
none
11 mcast-nc                 5       50000000  5       0      low
none

```

show interfaces terse

```

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

```


show interfaces terse (QFabric Systems)

```

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

```


show interfaces (GRE)

| | |
|---------------------------------|---|
| Syntax | <pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | <p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display status information about the specified generic routing encapsulation (GRE) interface. |
| Options | <p><i>interface-type</i>—On M Series and T Series routers and EX Series switches, the interface type is <i>gr-fpc/pic/port</i>.</p> <p>brief detail extensive terse—(Optional) Display the specified output level of interface information.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| | <div>  <p>NOTE: You can configure generic routing encapsulation (GRE) interfaces (<i>gre-x/y/z</i>) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the <i>MPLS Applications Feature Guide for Routing Devices</i> and the <i>Junos OS, Release 15.1</i>.</p> </div> |
| Required Privilege Level | view |
| List of Sample Output | <p>show interfaces (GRE) on page 360</p> <p>show interfaces brief (GRE) on page 360</p> <p>show interfaces detail (GRE) on page 360</p> <p>show interfaces (Layer 2 Services Over GRE Interfaces) on page 361</p> <p>show interfaces extensive (Layer 2 Services Over GRE Interfaces) on page 361</p> <p>show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 362</p> <p>show interfaces extensive (GRE) on page 363</p> |

Output Fields Table 37 on page 356 lists the output fields for the **show interfaces** (GRE) command. Output fields are listed in the approximate order in which they appear.

Table 37: GRE show interfaces Output Fields

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

Table 37: GRE show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|------------------------------------|---|------------------------------|
| Logical interface | Name of the logical interface. | All levels |
| Index | Logical interface index number, which reflects its initialization sequence. | detail extensive none |
| SNMP ifIndex | Logical interface SNMP interface index number. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support. | detail extensive |
| Flags | <p>Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under <i>Common Output Fields Description</i>. describe general information about the logical interface.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> • Reassemble-Pkts—If the Flags field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation. | All levels |
| IP-Header | <p>IP header of the logical interface. If the tunnel key statement is configured, this information is included in the IP Header entry.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> • df—If the IP-Header field includes this string immediately following the 16 bits of identification information (that is, if :df displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation. | All levels |
| Encapsulation | Encapsulation on the logical interface. | All levels |
| L2 Routing Instance | Name of the Layer 2 routing instance associated with the GRE interface. | All levels |
| L3 Routing Instance | Name of the Layer 3 routing instance associated with the GRE interface. | All levels |
| Copy-tos-to-outer-ip-header | <p>Status of type of service (ToS) bits in the GRE packet header:</p> <ul style="list-style-type: none"> • On—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel. • Off—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header. <p>NOTE: EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always Off in switch output.</p> | detail extensive |
| Gre keepalives configured | <p>Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel.</p> <p>NOTE: EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always Off in switch output.</p> | detail extensive |

Table 37: GRE show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|---------------------------------------|---|------------------------------|
| Gre keepalives adjacency state | Status of the other end of the GRE tunnel: Up or Down . If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up. | detail extensive |
| Input packets | Number of packets received on the logical interface. | None specified |
| Output packets | Number of packets transmitted on the logical interface. | None specified |
| Traffic statistics | <p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input rate—Rate of bits and packets received on the interface. • Output rate—Rate of bits and packets transmitted on the interface. | detail extensive |
| Local statistics | Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive |
| Transit statistics | Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. | detail extensive none |
| Protocol | Protocol family configured on the logical interface, such as iso , inet6 , or mpls . | detail extensive none |
| <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 |
| MTU | MTU size on the logical interface. | detail extensive none |
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |
| Route table | Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0 . | detail extensive |
| Flags | Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> . | detail extensive none |
| Addresses, Flags | Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . | 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 |

Table 37: GRE show interfaces Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------|---|-------------------------|
| Generation | Unique number for use by Juniper Networks technical support only. | detail extensive |

Sample Output

show interfaces (GRE)

```

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

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

```

show interfaces brief (GRE)

```

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

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

```

show interfaces detail (GRE)

```

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

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

```

```

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

```

show interfaces (Layer 2 Services Over GRE Interfaces)

```

user@host> show interfaces gr-2/2/10
show interfaces gr-2/2/10
Physical interface: gr-2/2/10, Enabled, Physical link is Up
Interface index: 214, SNMP ifIndex: 690
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

Logical interface gr-2/2/10.0 (Index 342) (SNMP ifIndex 10834)
Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
3.0.0.1:3.0.0.254:47:df:64:0000000000000000 Encapsulation: GRE-NULL
L2 Routing Instance: vs1, L3 Routing Instance: default
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Input packets : 2
Output packets: 0
Protocol bridge, MTU: 1476
Flags: Sendbroadcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6/8, Local: 6.0.0.1, Broadcast: 6.255.255.255

```

show interfaces extensive (Layer 2 Services Over GRE Interfaces)

```

user@host> show interfaces gr-2/2/10.0 extensive

Flags: SNMP-Traps Encapsulation: ENET2
L2 Routing Instance: vs1, L3 Routing Instance: default
Traffic statistics:
Input bytes :          58851250
Output bytes :           0
Input packets:        1279375
Output packets:         0
Local statistics:
Input bytes :           0
Output bytes :           0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :          58851250          75136 bps

```

```

Output bytes : 0 0 bps
Input packets: 1279375 204 pps
Output packets: 0 0 pps
Protocol bridge, MTU: 1476, Generation: 175, Route table: 7
Flags: Access-Mode

```

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

```

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

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 180.20.30.2:180.20.20.3:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
Input bytes : 5600886
Output bytes : 2881784
Input packets: 20010
Output packets: 10018
Local statistics:
Input bytes : 398
Output bytes : 264
Input packets: 5
Output packets: 3
Transit statistics:
Input bytes : 5600488 0 bps
Output bytes : 2881520 0 bps
Input packets: 20005 0 pps
Output packets: 10015 0 pps
Protocol inet, Generation: 159, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 90.90.90/24, Local: 90.90.90.10, Broadcast: 90.90.90.255,
Generation: 144

Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 160.20.40.2:160.20.30.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL

```



```

Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
  Input bytes :          260
  Output bytes :        2880148
  Input packets:          4
  Output packets:       10002
Local statistics:
  Input bytes :          112
  Output bytes :           0
  Input packets:         2
  Output packets:        0
Transit statistics:
  Input bytes :          148
  Output bytes :        2880148
  Input packets:         2
  Output packets:       10002
Protocol inet, Generation: 171, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 70.70.70/24, Local: 70.70.70.10, Broadcast: 70.70.70.255,
  Generation: 160

```

show interfaces extensive (GRE)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 360](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 362](#).

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 368</p> <p>show interfaces irb snmp-index on page 369</p> |
| Output Fields | Table 38 on page 364 lists the output fields for the show interfaces irb command. Output fields are listed in the approximate order in which they appear. |

Table 38: show interfaces irb Output Fields

| Field Name | Field Description | Level of Output |
|---------------------------|---|-----------------|
| Physical Interface | | |
| 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 |

Table 38: show interfaces irb Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------------|
| 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 |
| 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 |

Table 38: show interfaces irb 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. | 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 |
| 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

Table 38: show interfaces irb 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 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 | Speed at which the interface is running. | detail extensive |
| Routing Instance | Routing instance IRB is configured under. | detail extensive |
| Bridging Domain | Bridging domain IRB is participating in. | detail extensive |
| Traffic statistics | Number and rate of bytes and packets received and transmitted on the logical interface. <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 | 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 |

Table 38: show interfaces irb Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|-------------------------|---|-------------------------|
| 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, Runts: 0, Giants: 0, Policed discards:
0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface irb.0 (Index 68) (SNMP ifIndex 70) (Generation 143)
  Flags: Hardware-Down SNMP-Traps 0x4000 Encapsulation: ENET2
  Bandwidth: 1000mbps
  Routing Instance: customer_0 Bridging Domain: bd0

```

```

Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Protocol inet, MTU: 1500, Generation: 154, Route table: 0
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
    Destination: 10.51.1/24, 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: 10.51.1/24, Local: 10.51.1.2, Broadcast: 10.51.1.255
  Protocol multiservice, MTU: 1500
    Flags: Is-Primary

```

show interfaces queue

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

Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 39 on page 371](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 39 on page 371](#).

Table 39: Layer 2 Overhead and Transmitted Packets or Byte Counts

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

Layer 2 Statistics—Fragmentation Overhead Calculation

MLPPP/MC-MLPPP Overhead details:

=====

Fragment 1:

```

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

```

Fragments 2 .. n :

```

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

```

MLFR (FRF15) Overhead details:

=====

Fragment 1:

```

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

```

Fragments 2 ...n :

```

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

```

MFR (FRF16) Overhead details:

=====

```

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

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

```

Overhead with LFI

```

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

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

```

The following examples show overhead for different cases:

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

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

Additional Information

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



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

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

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

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

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

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

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

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

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

| | |
|---------------------------------|---|
| Required Privilege Level | view |
| List of Sample Output | show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC) on page 378 show interfaces queue (Aggregated Ethernet on a T320 Router) on page 379 show interfaces queue (Gigabit Ethernet on a T640 Router) on page 381 show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC) on page 381 show interfaces queue (Gigabit Ethernet IQ2 PIC) on page 385 show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC) on page 388 show interfaces queue ingress (Gigabit Ethernet IQ2 PIC) on page 390 show interfaces queue egress (Gigabit Ethernet IQ2 PIC) on page 391 show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC) on page 393 show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode) on page 395 show interfaces queue (QFX Series) on page 405 show interfaces queue l2-statistics (lsq interface) on page 406 show interfaces queue lsq (lsq-ifd) on page 407 |
| Output Fields | Table 40 on page 374 lists the output fields for the show interfaces queue command. Output fields are listed in the approximate order in which they appear. |

Table 40: show interfaces queue Output Fields

| Field Name | Field Description |
|---|---|
| Physical interface | Name of the physical interface. |
| Enabled | State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> . |
| Interface index | Physical interface's index number, which reflects its initialization sequence. |
| SNMP ifIndex | SNMP index number for the interface. |
| Forwarding classes supported | Total number of forwarding classes supported on the specified interface. |
| Forwarding classes in use | Total number of forwarding classes in use on the specified interface. |
| Ingress queues supported | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface. |
| Ingress queues in use | On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface. |
| Output queues supported | Total number of output queues supported on the specified interface. |
| Output queues in use | Total number of output queues in use on the specified interface. |
| Egress queues supported | Total number of egress queues supported on the specified interface. |
| Egress queues in use | Total number of egress queues in use on the specified interface. |
| Queue counters (Ingress) | CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <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. |
| Burst size | (Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface. |
| The following output fields are applicable to both interface component and Packet Forwarding component in the show interfaces queue command: | |
| Queue | Queue number. |
| Forwarding classes | Forwarding class name. |

Table 40: show interfaces queue Output Fields (*continued*)

| Field Name | Field Description |
|-----------------------------|--|
| Queued Packets | <p>Number of packets queued to this queue.</p> <p>NOTE: For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see “Additional Information” on page 372.</p> |
| Queued Bytes | <p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see Table 41 on page 377.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see “Additional Information” on page 372.</p> |
| Transmitted Packets | <p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.</p> <p>NOTE: For Layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 370</p> |
| Transmitted Bytes | <p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see Table 41 on page 377.</p> <p>NOTE: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p>NOTE: For Layer 2 statistics, see “Overhead for Layer 2 Statistics” on page 370</p> |
| Tail-dropped packets | <p>Number of packets dropped because of tail drop.</p> <p>NOTE: The Tail-dropped packets counter is not supported on the PTX Series Packet Transport Routers.</p> |
| RL-dropped packets | <p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see “Additional Information” on page 372.</p> <p>NOTE: The RL-dropped packets counter is not supported on the PTX Series Packet Transport Routers, and is omitted from the output.</p> |
| RL-dropped bytes | <p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs, MPCs, and Enhanced Queuing DPCs only, this statistic is not included in the queued traffic statistics. For more information, see “Additional Information” on page 372.</p> |

Table 40: show interfaces queue Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|--|
| RED-dropped packets | <p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. Low, TCP—Number of low-loss priority TCP packets dropped because of RED. High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. High, TCP—Number of high-loss priority TCP packets dropped because of RED. (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> Low—Number of low-loss priority packets dropped because of RED. Medium-low—Number of medium-low loss priority packets dropped because of RED. Medium-high—Number of medium-high loss priority packets dropped because of RED. High—Number of high-loss priority packets dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |
| RED-dropped bytes | <p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see Table 41 on page 377.</p> <ul style="list-style-type: none"> (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. High, TCP—Number of high-loss priority TCP bytes dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> |

Byte counts vary by interface hardware. [Table 41 on page 377](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware. [Table 41 on page 377](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 41: Byte Count by Interface Hardware

| Interface Hardware | Output Level | Byte Count Includes | Comments |
|----------------------------------|-----------------------------|---|--|
| Gigabit Ethernet IQ and IQE PICs | Interface | <p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p> | <p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p> |
| | Packet forwarding component | <p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> | — |
| Non-IQ PIC | Interface | <p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap. <p>M Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> Queued: The sum of the transmitted bytes and the RED dropped bytes. Transmitted: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble. RED dropped: Full Layer 2 overhead (including all L2 encapsulation and CRC) + 12 inter-packet gap + 8 for the preamble (does not include the VLAN header or MPLS pushed bytes). | <p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p> |

Table 41: Byte Count by Interface Hardware (*continued*)

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

Sample Output

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

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

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

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

show interfaces queue (Aggregated Ethernet on a T320 Router)

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

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

```

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

```

```

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

```

show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 42
  Forwarding classes: 8 supported, 8 in use
  Output queues: 8 supported, 8 in use
  Queue: 0, Forwarding classes: be
    Queued:
      Packets      :          13          0 pps
      Bytes        :         622          0 bps
    Transmitted:
      Packets      :          13          0 pps
      Bytes        :         622          0 bps
      Tail-dropped packets :          0          0 pps
      RED-dropped packets :          0          0 pps
      RED-dropped bytes  :          0          0 bps
  Queue: 1, Forwarding classes: af1
    Queued:
      Packets      :      1725947945      372178 pps
      Bytes        :      220921336960      381110432 bps
    Transmitted:
      Packets      :      1725947945      372178 pps
      Bytes        :      220921336960      381110432 bps
      Tail-dropped packets :          0          0 pps
      RED-dropped packets :          0          0 pps
      RED-dropped bytes  :          0          0 bps
  Queue: 2, Forwarding classes: ef1
    Queued:
      Packets      :          0          0 pps
      Bytes        :          0          0 bps
    Transmitted:
      Packets      :          0          0 pps
      Bytes        :          0          0 bps
      Tail-dropped packets :          0          0 pps
      RED-dropped packets :          0          0 pps
      RED-dropped bytes  :          0          0 bps
  Queue: 3, Forwarding classes: nc
    Queued:
      Packets      :          571          0 pps
      Bytes        :         49318         336 bps
    Transmitted:
      Packets      :          571          0 pps
      Bytes        :         49318         336 bps
      Tail-dropped packets :          0          0 pps
      RED-dropped packets :          0          0 pps
      RED-dropped bytes  :          0          0 bps

```

show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate

```

```

Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
  Forwarding classes: 16 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
  Queue: 0, Forwarding classes: best-effort
    Queued:
      Packets      :      148450735      947295 pps
      Bytes        :      8016344944    409228848 bps
    Transmitted:
      Packets      :      76397439      487512 pps
      Bytes        :    4125461868    210602376 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :      72053285      459783 pps
        Low        :      72053285      459783 pps
        Medium-low  :           0          0 pps
        Medium-high :           0          0 pps
        High        :           0          0 pps
      RED-dropped bytes :    3890877444    198626472 bps
        Low        :    3890877444    198626472 bps
        Medium-low  :           0          0 bps
        Medium-high :           0          0 bps
        High        :           0          0 bps
  Queue: 1, Forwarding classes: expedited-forwarding
    Queued:
      Packets      :           0          0 pps
      Bytes        :           0          0 bps
    Transmitted:
      Packets      :           0          0 pps
      Bytes        :           0          0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :           0          0 pps
        Low        :           0          0 pps
        Medium-low  :           0          0 pps
        Medium-high :           0          0 pps
        High        :           0          0 pps
      RED-dropped bytes :           0          0 bps
        Low        :           0          0 bps
        Medium-low  :           0          0 bps
        Medium-high :           0          0 bps
        High        :           0          0 bps
  Queue: 2, Forwarding classes: assured-forwarding
    Queued:
      Packets      :      410278257      473940 pps
      Bytes        :    22156199518    204742296 bps
    Transmitted:
      Packets      :      4850003      4033 pps
      Bytes        :    261900162    1742256 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :      405425693      469907 pps
        Low        :      405425693      469907 pps
        Medium-low  :           0          0 pps
        Medium-high :           0          0 pps
        High        :           0          0 pps
      RED-dropped bytes :    21892988124    203000040 bps
        Low        :    21892988124    203000040 bps
        Medium-low  :           0          0 bps
        Medium-high :           0          0 bps
        High        :           0          0 bps
  Queue: 3, Forwarding classes: network-control
    Queued:
      Packets      :           0          0 pps

```

```

Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : 76605230 485376 pps
Bytes : 5209211400 264044560 bps
Transmitted:
Packets : 76444631 484336 pps
Bytes : 5198235612 263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets : 160475 1040 pps
Low : 160475 1040 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 10912300 565760 bps
Low : 10912300 565760 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low : 0 0 pps
Medium-low : 0 0 pps
Medium-high : 0 0 pps
High : 0 0 pps
RED-dropped bytes : 0 0 bps
Low : 0 0 bps
Medium-low : 0 0 bps
Medium-high : 0 0 bps
High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 4836136 3912 pps
Bytes : 333402032 2139056 bps
Transmitted:
Packets : 3600866 1459 pps
Bytes : 244858888 793696 bps
Tail-dropped packets : Not Available

```

```

RED-dropped packets :          1225034          2450 pps
  Low                :          1225034          2450 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :          83302312        1333072 bps
  Low                :          83302312        1333072 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Transmitted:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Tail-dropped packets : Not Available
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :              0              0 bps
  Low                :              0              0 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps

```

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
  Packets            :          77059796        486384 pps
  Bytes              :          3544750624      178989576 bps
Transmitted:
  Packets            :          77059797        486381 pps
  Bytes              :          3544750670      178988248 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps
  High               :              0              0 pps
RED-dropped bytes   :              0              0 bps
  Low                :              0              0 bps
  Medium-low         :              0              0 bps
  Medium-high        :              0              0 bps
  High               :              0              0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Transmitted:
  Packets            :              0              0 pps
  Bytes              :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets :              0              0 pps
  Low                :              0              0 pps
  Medium-low         :              0              0 pps
  Medium-high        :              0              0 pps

```

```

      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 4846580 3934 pps
    Bytes : 222942680 1447768 bps
  Transmitted:
    Packets : 4846580 3934 pps
    Bytes : 222942680 1447768 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps

```

show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

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

```

```

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

```



```

Tail-dropped packets : Not Available
RL-dropped packets   :                0          0 pps
RL-dropped bytes     :                0          0 bps
RED-dropped packets   :                0          0 pps
RED-dropped bytes     :                0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets              :                77009          11 pps
Bytes                :            6894286          7888 bps
Transmitted:
Packets              :                77009          11 pps
Bytes                :            6894286          7888 bps
Tail-dropped packets : Not Available
RL-dropped packets   :                0          0 pps
RL-dropped bytes     :                0          0 bps
RED-dropped packets   :                0          0 pps
RED-dropped bytes     :                0          0 bps

```

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

```

Queued:
Packets              :                1031          0 pps
Bytes                :            147328          0 bps
Transmitted:
Packets              :                1031          0 pps
Bytes                :            147328          0 bps
Tail-dropped packets :                0          0 pps
RED-dropped packets   :                0          0 pps
Low, non-TCP         :                0          0 pps
Low, TCP              :                0          0 pps
High, non-TCP        :                0          0 pps
High, TCP             :                0          0 pps
RED-dropped bytes     :                0          0 bps
Low, non-TCP         :                0          0 bps
Low, TCP              :                0          0 bps
High, non-TCP        :                0          0 bps
High, TCP             :                0          0 bps

```

Queue: 1, Forwarding classes: expedited-forwarding

```

Queued:
Packets              :                0          0 pps
Bytes                :                0          0 bps
Transmitted:
Packets              :                0          0 pps
Bytes                :                0          0 bps
Tail-dropped packets :                0          0 pps
RED-dropped packets   :                0          0 pps
Low, non-TCP         :                0          0 pps
Low, TCP              :                0          0 pps
High, non-TCP        :                0          0 pps
High, TCP             :                0          0 pps
RED-dropped bytes     :                0          0 bps
Low, non-TCP         :                0          0 bps
Low, TCP              :                0          0 bps
High, non-TCP        :                0          0 bps
High, TCP             :                0          0 bps

```

Queue: 2, Forwarding classes: assured-forwarding

```

Queued:
Packets              :                0          0 pps
Bytes                :                0          0 bps
Transmitted:

```

```

Packets          : 0 0 pps
Bytes            : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
  Low, non-TCP    : 0 0 pps
  Low, TCP        : 0 0 pps
  High, non-TCP   : 0 0 pps
  High, TCP       : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low, non-TCP    : 0 0 bps
  Low, TCP        : 0 0 bps
  High, non-TCP   : 0 0 bps
  High, TCP       : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets          : 94386 12 pps
  Bytes            : 13756799 9568 bps
Transmitted:
  Packets          : 94386 12 pps
  Bytes            : 13756799 9568 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
  Low, non-TCP    : 0 0 pps
  Low, TCP        : 0 0 pps
  High, non-TCP   : 0 0 pps
  High, TCP       : 0 0 pps
RED-dropped bytes : 0 0 bps
  Low, non-TCP    : 0 0 bps
  Low, TCP        : 0 0 bps
  High, non-TCP   : 0 0 bps
  High, TCP       : 0 0 bps

```

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

```

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

```

```

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

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
  Transmitted:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
  Transmitted:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :      9397      0 pps
    Bytes        :      3809052      232 bps
  Transmitted:
    Packets      :      9397      0 pps
    Bytes        :      3809052      232 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

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

```

```

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

```

show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

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

```

```

Packets          : Not Available
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets          : Not Available
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets          :                      80564692      0 pps
Bytes            :                      3383717100     0 bps
Transmitted:
Packets          :                      80564692      0 pps
Bytes            :                      3383717100     0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets          :                      80564685      0 pps
Bytes            :                      3383716770     0 bps
Transmitted:
Packets          :                      80564685      0 pps
Bytes            :                      3383716770     0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Transmitted:
Packets          :                      0          0 pps
Bytes            :                      0          0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets          :                      9538          0 pps
Bytes            :                      3819840        0 bps
Transmitted:
Packets          :                      9538          0 pps
Bytes            :                      3819840        0 bps
Tail-dropped packets :                      0          0 pps
RED-dropped packets :                      0          0 pps
RED-dropped bytes  :                      0          0 bps

```

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

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :          110208969          472875 pps
    Bytes        :          5951284434        204282000 bps
  Transmitted:
    Packets      :          110208969          472875 pps
    Bytes        :          5951284434        204282000 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps
    High         :              0              0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Transmitted:
    Packets      :              0              0 pps
    Bytes        :              0              0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets :
    Low          :              0              0 pps
    Medium-low   :              0              0 pps
    Medium-high  :              0              0 pps
    High         :              0              0 pps
  RED-dropped bytes :
    Low          :              0              0 bps
    Medium-low   :              0              0 bps
    Medium-high  :              0              0 bps

```

```

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

```



```

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

```

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

```

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

  Interface index: 192, SNMP ifIndex: 1948

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

  Forwarding classes: 16 supported, 9 in use

  Egress queues: 8 supported, 8 in use

  Queue: 0, Forwarding classes: DEFAULT

  Queued:

    Packets          :                214886                13449 pps

    Bytes            :                9884756                5164536 bps

  Transmitted:

    Packets          :                214886                13449 pps

    Bytes            :                9884756                5164536 bps

```

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

Queue: 1, Forwarding classes: REALTIME

Queued:

| | | |
|-----------|---|-------|
| Packets : | 0 | 0 pps |
| Bytes : | 0 | 0 bps |

Transmitted:

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

Queue: 2, Forwarding classes: PRIVATE

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

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

Queue: 3, Forwarding classes: CONTROL

Queued:

| | | | |
|---------|---|------|-------|
| Packets | : | 60 | 0 pps |
| Bytes | : | 4560 | 0 bps |

Transmitted:

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

| | | | |
|-------------|---|---|-------|
| Low | : | 0 | 0 bps |
| Medium-low | : | 0 | 0 bps |
| Medium-high | : | 0 | 0 bps |
| High | : | 0 | 0 bps |

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

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

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

| | | | |
|----------------------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |

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

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

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

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
|---------|---|---|-------|

```

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

```

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

```

Queued:
Packets        :            371365            23620 pps
Bytes          :          15597330          7936368 bps
Transmitted:
Packets        :            371365            23620 pps
Bytes          :          15597330          7936368 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets :                0                0 pps
Low            :                0                0 pps
Medium-low    :                0                0 pps
Medium-high   :                0                0 pps

```

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

Queue: 1, Forwarding classes: REALTIME

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

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

Queue: 2, Forwarding classes: PRIVATE

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

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

Queue: 3, Forwarding classes: CONTROL

Queued:

| | | |
|-----------|---------|--------|
| Packets : | 32843 | 0 pps |
| Bytes : | 2641754 | 56 bps |

Transmitted:

| | | |
|------------------------|---------|--------|
| Packets : | 32843 | 0 pps |
| Bytes : | 2641754 | 56 bps |
| Tail-dropped packets : | 0 | 0 pps |
| RED-dropped packets : | 0 | 0 pps |
| Low : | 0 | 0 pps |
| Medium-low : | 0 | 0 pps |
| Medium-high : | 0 | 0 pps |
| High : | 0 | 0 pps |
| RED-dropped bytes : | 0 | 0 bps |
| Low : | 0 | 0 bps |
| Medium-low : | 0 | 0 bps |
| Medium-high : | 0 | 0 bps |
| High : | 0 | 0 bps |

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

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

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

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

| | | | |
|-------------|---|---|-------|
| Low | : | 0 | 0 bps |
| Medium-low | : | 0 | 0 bps |
| Medium-high | : | 0 | 0 bps |
| High | : | 0 | 0 bps |

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

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

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

| | | | |
|---------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |

Transmitted:

| | | | |
|----------------------|---|---|-------|
| Packets | : | 0 | 0 pps |
| Bytes | : | 0 | 0 bps |
| Tail-dropped packets | : | 0 | 0 pps |

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

show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0      0 pps
    Bytes       : 0      0 bps
  Transmitted:
    Packets      : 0      0 pps
    Bytes       : 0      0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0      0 pps
    Total-dropped bytes  : 0      0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0      0 pps
    Bytes       : 0      0 bps
  Transmitted:
    Packets      : 0      0 pps
    Bytes       : 0      0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0      0 pps
    Total-dropped bytes  : 0      0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0      0 pps
    Bytes       : 0      0 bps
  Transmitted:
    Packets      : 0      0 pps
    Bytes       : 0      0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0      0 pps
    Total-dropped bytes  : 0      0 bps
Queue: 7, Forwarding classes: network-control

```

```

Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available
Total-dropped packets:      0      0 pps
Total-dropped bytes  :      0      0 bps
Queue: 8, Forwarding classes: mcast
Queued:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Transmitted:
  Packets      :      0      0 pps
  Bytes       :      0      0 bps
Tail-dropped packets : Not Available
Total-dropped packets:      0      0 pps
Total-dropped bytes  :      0      0 bps

```

show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
Queued:
  Packets      :      1      0 pps
  Bytes       :    1001      0 bps
Transmitted:
  Packets      :      5      0 pps
  Bytes       :    1062      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets  :      0      0 pps
RED-dropped bytes    :      0      0 bps
Queue: 1, Forwarding classes: ef
Queued:
  Packets      :      1      0 pps
  Bytes       :    1500      0 bps
Transmitted:
  Packets      :      6      0 pps
  Bytes       :    1573      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets  :      0      0 pps
RED-dropped bytes    :      0      0 bps
Queue: 2, Forwarding classes: af
Queued:
  Packets      :      1      0 pps
  Bytes       :     512      0 bps
Transmitted:
  Packets      :      3      0 pps
  Bytes       :     549      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets  :      0      0 pps
RED-dropped bytes    :      0      0 bps
Queue: 3, Forwarding classes: nc
Queued:
  Packets      :      0      0 pps

```

```

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

```

show interfaces queue lsq (lsq-ifd)

```

user@switch> show interfaces queue lsq-1/0/0
Logical interface lsq-1/0/0 (Index 348) (SNMP ifIndex 660)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets          :          55576          1206 pps
    Bytes            :      29622008      5145472 bps
  Transmitted:
    Packets          :          55576          1206 pps
    Bytes            :      29622008      5145472 bps
    Tail-dropped packets :                0                0 pps
    RL-dropped packets  :                0                0 pps
    RL-dropped bytes    :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low                :                0                0 pps
    Medium-low         :                0                0 pps
    Medium-high        :                0                0 pps
    High               :                0                0 pps
    RED-dropped bytes   :                0                0 bps
    Low                :                0                0 bps
    Medium-low         :                0                0 bps
    Medium-high        :                0                0 bps
    High               :                0                0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
  Transmitted:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps
    Tail-dropped packets :                0                0 pps
    RL-dropped packets  :                0                0 pps
    RL-dropped bytes    :                0                0 bps
    RED-dropped packets :                0                0 pps
    Low                :                0                0 pps
    Medium-low         :                0                0 pps
    Medium-high        :                0                0 pps
    High               :                0                0 pps
    RED-dropped bytes   :                0                0 bps
    Low                :                0                0 bps
    Medium-low         :                0                0 bps
    Medium-high        :                0                0 bps
    High               :                0                0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets          :                0                0 pps
    Bytes            :                0                0 bps

```

```

Transmitted:
  Packets          :          0          0 pps
  Bytes            :          0          0 bps
  Tail-dropped packets :          0          0 pps
  RL-dropped packets :          0          0 pps
  RL-dropped bytes  :          0          0 bps
  RED-dropped packets :          0          0 pps
    Low            :          0          0 pps
    Medium-low     :          0          0 pps
    Medium-high    :          0          0 pps
    High           :          0          0 pps
  RED-dropped bytes :          0          0 bps
    Low            :          0          0 bps
    Medium-low     :          0          0 bps
    Medium-high    :          0          0 bps
    High           :          0          0 bps
Queue: 3, Forwarding classes: nc
Queued:
  Packets          :        22231        482 pps
  Bytes            :       11849123     2057600 bps
Transmitted:
  Packets          :        22231        482 pps
  Bytes            :       11849123     2057600 bps
  Tail-dropped packets :          0          0 pps
  RL-dropped packets :          0          0 pps
  RL-dropped bytes  :          0          0 bps
  RED-dropped packets :          0          0 pps
    Low            :          0          0 pps
    Medium-low     :          0          0 pps
    Medium-high    :          0          0 pps
    High           :          0          0 pps
  RED-dropped bytes :          0          0 bps
    Low            :          0          0 bps
    Medium-low     :          0          0 bps
    Medium-high    :          0          0 bps
    High           :          0          0 bps

```

show interfaces xe

| | |
|---------------------------------|--|
| Syntax | <pre>show interfaces <i>device-name:type-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <routing-instance (all <i>instance-name</i>)> <snmp-index <i>snmp-index</i>> <statistics></pre> |
| Release Information | <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> |
| Description | Display status information about the specified 10-Gigabit Ethernet interface. This command does not display statistics for routed VLAN interfaces. |
| Options | <p><i>device-name:type-fpc/pic/port</i>—(QFabric systems only) The device name is either the serial number or the alias of the QFabric system component, such as a Node device, Interconnect device, or QFabric infrastructure. The name must contain a maximum of 128 characters and not contain any colons.</p> <p>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 (all <i>instance-name</i>)—(Optional) Display the name of an individual routing instance or display all routing instances.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Monitoring Interface Status and Traffic on page 79 • Troubleshooting Network Interfaces on page 80 • Troubleshooting an Aggregated Ethernet Interface on page 146 • Junos OS Network Interfaces Library for Routing Devices |
| List of Sample Output | <p>show interfaces on page 417</p> <p>show interfaces (Asymmetric Flow Control) on page 418</p> <p>show interfaces brief on page 418</p> <p>show interfaces detail on page 418</p> <p>show interfaces detail (Asymmetric Flow Control) on page 420</p> <p>show interfaces extensive on page 421</p> |

[show interfaces extensive \(Asymmetric Flow Control\) on page 423](#)

[show interfaces terse on page 425](#)

[show interfaces \(QFabric System\) on page 425](#)

Output Fields Table 42 on page 410 lists the output fields for the **show interfaces xe** command. Output fields are listed in the approximate order in which they appear.

Table 42: 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 42: show interfaces xe Output Fields (*continued*)

| Field Name | Field Description | Level of Output |
|--------------------------------|--|------------------------------|
| Configured-flow-control | <p>Configured flow control for the interface transmit buffers (tx-buffers) and receive buffers (rx-buffers):</p> <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 | <p>Remote fault status:</p> <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 42: 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 42: 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 42: 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 42: 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 42: 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 42: 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 interlace. | 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:

| | Receive | Transmit |
|--------------------|---------|----------|
| Total octets | 0 | 0 |
| Total packets | 0 | 0 |
| Unicast packets | 0 | 0 |
| Broadcast packets | 0 | 0 |
| Multicast packets | 0 | 0 |
| CRC/Align errors | 0 | 0 |
| FIFO errors | 0 | 0 |
| MAC control frames | 0 | 0 |
| MAC pause frames | 0 | 0 |
| Oversized frames | 0 | |
| Jabber frames | 0 | |
| Fragment frames | 0 | |
| VLAN tagged frames | 0 | |
| Code violations | 0 | |

MAC Priority Flow Control Statistics:

| | | |
|--------------|---|---|
| Priority : 0 | 0 | 0 |
| Priority : 1 | 0 | 0 |
| Priority : 2 | 0 | 0 |
| Priority : 3 | 0 | 0 |
| Priority : 4 | 0 | 0 |
| Priority : 5 | 0 | 0 |
| Priority : 6 | 0 | 0 |
| Priority : 7 | 0 | 0 |

Filter statistics:

| | | |
|----------------------|---|---|
| Input packet count | 0 | |
| Input packet rejects | 0 | |
| Input DA rejects | 0 | |
| Input SA rejects | 0 | |
| Output packet count | | 0 |

```

Output packet pad count                                0
Output packet error count                              0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      usec
0 best-effort              75      7500000000  75      0      low
none
7 network-control          5       500000000   5       0      low
none
8 mcast                    20      2000000000  20      0      low
none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Local statistics:
  Input bytes : 0
  Output bytes : 0
  Input packets: 0
  Output packets: 0
Transit statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces extensive (Asymmetric Flow Control)

```

user@switch> show interfaces xe-0/0/1 extensive
Physical interface: xe-0/0/1, Enabled, Physical link is Up
Interface index: 49195, SNMP ifIndex: 591, Generation: 169
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled,
Configured-flow-control tx-buffers: off rx-buffers: on
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 12 supported, 12 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1d:b5:f7:4e:e1, Hardware address: 00:1d:b5:f7:4e:e1
Last flapped : 2011-06-01 00:42:03 PDT (00:03:08 ago)
Statistics last cleared: 2011-06-01 00:44:39 PDT (00:00:32 ago)
Traffic statistics:
  Input bytes : 0 0 bps
  Output bytes : 0 0 bps
  Input packets: 0 0 pps
  Output packets: 0 0 pps
IPv6 transit statistics:
  Input bytes : 0

```

```

Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 12 supported, 9 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          0              0              0
1 fc7                  0              0              0
2 no-loss              0              0              0
3 fcoe                 0              0              0
4 fc4                  0              0              0
5 fc5                  0              0              0
6 fc6                  0              0              0
7 network-cont         0              0              0
8 mcast                0              0              0

Queue number:      Mapped forwarding classes
0                  best-effort
1                  fc7
2                  no-loss
3                  fcoe
4                  fc4
5                  fc5
6                  fc6
7                  network-control
8                  mcast

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets    0            0
Unicast packets  0            0
Broadcast packets 0            0
Multicast packets 0            0
CRC/Align errors 0            0
FIFO errors       0            0
MAC control frames 0            0
MAC pause frames  0            0
Oversized frames  0
Jabber frames     0
Fragment frames   0
VLAN tagged frames 0
Code violations    0
MAC Priority Flow Control Statistics:
Priority : 0       0            0
Priority : 1       0            0

```

```

Priority : 2          0          0
Priority : 3          0          0
Priority : 4          0          0
Priority : 5          0          0
Priority : 6          0          0
Priority : 7          0          0
Filter statistics:
Input packet count    0
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
Direction : Output
CoS transmit queue    Bandwidth      Buffer Priority  Limit
                        %      bps      %      usec
0 best-effort         75    7500000000    75      0    low    none
7 network-control     5     500000000     5      0    low    none
8 mcast               20    2000000000    20      0    low    none

Logical interface xe-0/0/1.0 (Index 73) (SNMP ifIndex 523) (Generation 143)
Flags: SNMP-Traps 0x0 Encapsulation: ENET2
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0          0 bps
Output bytes : 0          0 bps
Input packets: 0          0 pps
Output packets: 0          0 pps
Protocol eth-switch, MTU: 0, Generation: 170, Route table: 0
Flags: Trunk-Mode

```

show interfaces terse

```

user@switch> show interfaces xe-0/0/1 terse
Interface      Admin Link Proto  Local      Remote

xe-0/0/1       up    up
xe-0/0/1.0     up    up    eth-switch

```

show interfaces (QFabric System)

```

user@switch> show interfaces node1:xe-0/0/0
Physical interface: node1:xe-0/0/0, Enabled, Physical link is Down
Interface index: 129, SNMP ifIndex: 2884086
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex, BPDU
Error: None, MAC-REWRITE Error: None,
Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled

```

Interface flags: Internal: 0x4000
CoS queues : 8 supported, 8 maximum usable queues
Current address: 02:00:09:03:00:00, Hardware address: 02:00:09:03:00:00
Last flapped : Never
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

CHAPTER 22

LAGs and LACP Operational Commands

- `show lacp interfaces`
- `show lacp statistics interfaces` (View)

show lacp interfaces

| | |
|---------------------------------|---|
| Syntax | show lacp interfaces <interface-name> |
| Release Information | <p>Command introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p> <p>Command introduced in Junos OS Release 14.2R3</p> |
| Description | Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface. |
| Options | <p>none—Display LACP information for all interfaces.</p> <p><i>interface-name</i>—(Optional) Display LACP information for the specified interface:</p> <ul style="list-style-type: none"> • Aggregated Ethernet—<i>aex</i> • Gigabit Ethernet—<i>ge-fpc/pic/port</i> • 10-Gigabit Ethernet—<i>xe-fpc/pic/port</i> |
| Required Privilege Level | view |
| 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> • Example: Configuring Link Aggregation Between a QFX Series Product and an Aggregation Switch on page 136 • Configuring Aggregated Ethernet Links (CLI Procedure) • Configuring Link Aggregation on page 133 • Configuring Aggregated Ethernet LACP (CLI Procedure) • Configuring Aggregated Ethernet LACP on page 132 • Configuring LACP Link Protection of Aggregated Ethernet Interfaces (CLI Procedure) • Understanding Aggregated Ethernet Interfaces and LACP • Understanding Aggregated Ethernet Interfaces and LACP on page 129 • Junos OS Interfaces Fundamentals Configuration Guide |
| List of Sample Output | <p>show lacp interfaces (EX Series Switches) on page 430</p> <p>show lacp interfaces (QFX Series) on page 431</p> |

Output Fields Table 43 on page 429 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 43: show lacp interfaces Output Fields

| Field Name | Field Description |
|----------------------|--|
| Aggregated interface | Aggregated Ethernet interface name. |
| LACP State | <p>LACP state information for each aggregated Ethernet interface:</p> <ul style="list-style-type: none"> For a child interface configured with the force-up statement, LACP state displays FUP along with the interface name. Role—Role played by the interface. It can be one of the following: <ul style="list-style-type: none"> Actor—Local device participating in the LACP negotiation. Partner—Remote device participating in the LACP negotiation. Exp—Expired state. Yes indicates that the actor or partner is in an expired state. No indicates that the actor or partner is not in an expired state. Def—Default. Yes indicates that the actor's receive machine is using the default operational partner information, which is administratively configured for the partner. No indicates that the operational partner information in use has been received in an LACP PDU. Dist—Distribution of outgoing frames. No indicates that the distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is Yes. Col—Collection of incoming frames. Yes indicates that the collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is No. Syn—Synchronization. If the value is Yes, the link is considered to be synchronized. The link has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is No, the link is not synchronized. The link is currently not in the right aggregation. Aggr—Ability of the aggregation port to aggregate (Yes) or to operate only as an individual link (No). Timeout—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or a fast transmission rate, depending upon the expressed LACP timeout preference (Long Timeout or Short Timeout). Activity—Actor's or partner's port activity. Passive indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is Active. Active indicates the port's preference to participate in the protocol regardless of the partner's control value. |

Table 43: show lacp interfaces Output Fields (*continued*)

| Field Name | Field Description |
|---------------|--|
| LACP Protocol | <p>LACP protocol information for each aggregated interface:</p> <ul style="list-style-type: none"> 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 the state machine. The transmit state is 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—The multiplexer state machine initiates the process of attaching the port to the selected aggregator. Collecting—Yes indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. No indicates the receive function of this link is not enabled. Collecting distributing—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution. Detached—Process of detaching the port from the aggregator is in progress. Distributing—Yes indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames can be passed down from the aggregator's distribution function for transmission. No indicates the transmit function of this link is not enabled. Waiting—The multiplexer state machine is in a holding process, awaiting an outcome. |

Sample Output

show lacp interfaces (EX Series Switches)

```

user@switch> show lacp interfaces ae5
Aggregated interface: ae5
  LACP state:      Role  Exp  Def  Dist  Col  Syn  Aggr  Timeout  Activity
    xe-2/0/7      Actor  No   No   Yes   Yes  Yes   Yes    Fast    Active
    xe-2/0/7      Partner No   No   Yes   Yes  Yes   Yes    Fast    Passive

```

| | | | | | | | | | |
|----------|---------|----|----|----|-----|-----|-----|------|---------|
| xe-4/0/7 | Actor | No | No | No | No | No | Yes | Fast | Active |
| xe-4/0/7 | Partner | No | No | No | Yes | Yes | Yes | Fast | Passive |

| LACP protocol: | Receive State | Transmit State | Mux State |
|--------------------|---------------|----------------|-------------------------|
| xe-2/0/7(Active) | Current | Fast periodic | Collecting distributing |
| xe-34/0/7(Standby) | Current | Fast periodic | Waiting |

show lacp interfaces (QFX Series)

```
user@switch> show lacp interfaces nodegroup1:ae0 extensive
```

```
Aggregated interface: nodegroup1:ae0
```

| LACP state: | Role | Exp | Def | Dist | Col | Syn | Aggr | Timeout | Activity |
|-------------------|---------|-----|-----|------|-----|-----|------|---------|----------|
| node1:xe-0/0/1FUP | Actor | No | Yes | No | No | No | No | Yes | Fast |
| Active | | | | | | | | | |
| node1xe-0/0/1FUP | Partner | No | Yes | No | No | No | No | Yes | Fast |
| Passive | | | | | | | | | |
| node2:xe-0/0/2 | Actor | No | Yes | No | No | No | No | Yes | Fast |
| Active | | | | | | | | | |
| node2:xe-0/0/2 | Partner | No | Yes | No | No | No | No | Yes | Fast |
| Passive | | | | | | | | | |

| | | | |
|--------------------------|---------------|----------------|------------|
| LACP protocol: | Receive State | Transmit State | Mux State |
| node1:xe-0/0/1FUP | Current | Fast periodic | Collecting |
| distributing | | | |
| node2:xe-0/0/2 | Current | Fast periodic | Collecting |
| distributing | | | |
| node1:xe-0/0/1 (active) | Current | Fast periodic | Collecting |
| distributing | | | |
| node2:xe-0/0/2 (standby) | Current | Fast periodic | WAITING |

show lacp statistics interfaces (View)

| | |
|---------------------------------|---|
| Syntax | <code>show lacp statistics interfaces <i>interface-name</i></code> |
| Release Information | Command modified in Junos OS Release 10.2. |
| 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 |
| List of Sample Output | show lacp statistics interfaces on page 433 |
| Output Fields | Table 44 on page 433 lists the output fields for the <code>show lacp statistics interfaces</code> command. Output fields are listed in the approximate order in which they appear. |

Table 44: show lacp statistics interfaces Output Fields

| Field Name | Field Description |
|----------------------|--|
| Aggregated interface | Aggregated interface value. |
| LACP Statistics | <p>LACP statistics provide the following information:</p> <ul style="list-style-type: none"> 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. |

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 23

Redundant Trunk Group Operational Commands

- `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"> • Example: Configuring Redundant Trunk Links for Faster Recovery • Example: Configuring Redundant Trunk Links for Faster Recovery on page 153 • Understanding Redundant Trunk Links on page 151 |
| List of Sample Output | show redundant-trunk-group group-name Group1 on page 436 |
| Output Fields | Table 45 on page 436 lists the output fields for the show redundant-trunk-group command. Output fields are listed in the approximate order in which they appear. |

Table 45: 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 |

CHAPTER 24

Resilient Hashing Operational Commands

- `show forwarding-options enhanced-hash-key`

show forwarding-options enhanced-hash-key

| | |
|---------------------------------|--|
| Syntax | show forwarding-options enhanced-hash-key |
| Release Information | <p>Command introduced in Junos OS Release 13.2X51-D15 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D20 for QFX Series devices.</p> <p>Fabric Load Balancing Options output fields introduced in Junos OS Release 14.1X53-D10.</p> |
| Description | <p>Display information about which packet fields are used by the hashing algorithm to make hashing decisions.</p> <p>You can configure the fields that are inspected by the hashing algorithm to make hashing decisions for traffic entering a LAG bundle using the forwarding-options enhanced-hash-key statement.</p> |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic (CLI Procedure) on page 169 • Understanding the Algorithm Used to Hash LAG Bundle and Egress Next-Hop ECMP Traffic on page 161 • enhanced-hash-key on page 298 |
| List of Sample Output | <p>show forwarding-options enhanced-hash-key (Layer 2 Payload Hash Mode) on page 442</p> <p>show forwarding-options enhanced-hash-key (Layer 2 Header Hash Mode) on page 442</p> <p>show forwarding-options enhanced-hash-key (Fabric Load Balancing Options) on page 443</p> <p>show forwarding-options enhanced-hash-key (QFX10002 and QFX 10008 Switches) on page 443</p> |
| Output Fields | <p>Table 46 on page 440 lists the output fields for the show forwarding-options enhanced-hash-key command. Output fields are listed in the approximate order in which they first appear. Output fields vary by platform.</p> |

Table 46: show forwarding-options enhanced-hash-key Output Fields

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

Table 46: show forwarding-options enhanced-hash-key Output Fields (*continued*)

| Field Name | Field Description |
|--------------------------------|--|
| Destination IPv4 Addr | Indicates whether the Destination IPv4 Addr field is or is not used by the hashing algorithm: Yes or No. |
| Source IPv4 Addr | Indicates whether the Source IPv4 Addr field is or is not used by the hashing algorithm: Yes or No. |
| Vlan id | Indicates whether the Vlan ID field is or is not used by the hashing algorithm: Yes or No. |
| Inner-Vlan ID | Indicates whether the inner Vlan field is or is not used by the hashing algorithm: Yes or No. |
| Next Hdr | Indicates whether the Next Hdr field is or is not used by the hashing algorithm: Yes or No. |
| Destination IPv6 Addr | Indicates whether the Destination IPv6 Addr field is or is not used by the hashing algorithm: Yes or No. |
| Source IPv6 Addr | Indicates whether the Source IPv6 Addr field is or is not used by the hashing algorithm: Yes or No. |
| Ether Type | Indicates whether the Ether Type field is or is not used by the hashing algorithm: Yes or No. |
| Destination MAC Address | Indicates whether the Destination MAC Address field is or is not used by the hashing algorithm: Yes or No. |
| Source MAC Address | Indicates whether the Source MAC Address field is or is not used by the hashing algorithm: Yes or No. |
| Load Balancing Method | Indicates the load balancing method for adaptive load balancing (ALB): flowlet or per-packet. The load balancing method is flowlet by default, and can be configured using the fabric-load-balance statement. |
| Fabric Link Scale | Indicates the fabric link scale, in mbps. |
| Inactivity Interval | Indicates the fabric load balance inactivity interval, in microseconds (us). The inactivity interval is 16 microseconds by default, and can be configured using the inactivity-interval statement. |
| Hash Region Size/Trunk | Indicates the hash region size, in buckets per fabric trunk. |
| Seed | A hash seed value, between 0 and 4294967295. If a hash-seed value is not configured it is automatically assigned on the QFX10002 and QFX10008 switches. A hash-seed prevents traffic polarization to same links on the next hop QFX switch when two are connected with LAG/ECMP. |

Table 46: show forwarding-options enhanced-hash-key Output Fields (*continued*)

| Field Name | Field Description |
|---------------------|---|
| Key | Indicates whether the GRE key field is or is not used by the hashing algorithm: Yes or No. |
| Protocol | Indicates if a Generic Router Encapsulation (GRE) endpoint over routes was dynamically learned by a routing protocol such as RIP or OSPF. |
| MPLS Enabled | Indicates if MPLS is enabled under L2 switching. |
| VXLAN VNID | A 24-bit virtual network identifier (VNID) that uniquely identifies the Virtual Extensible Local Area Networks (VXLAN) segment. |

Sample Output

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

```

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

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

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

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

```

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

```

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

Current Hash Settings
-----

Hash-Mode                               : layer2-header

```


layer2 Hash settings-

layer2 packet fields

| | |
|-------------------------|-------|
| Ether Type | : Yes |
| Destination MAC Address | : Yes |
| Source MAC Address | : Yes |
| VLAN ID | : No |

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

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

Fabric Load Balancing Options

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

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

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

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

Inet settings:

| | |
|----------------------|-----|
| IPv4 dest address: | Yes |
| IPv4 source address: | Yes |
| L4 Dest Port: | Yes |
| L4 Source Port: | Yes |

Inet6 settings:

| | |
|----------------------|-----|
| IPv6 dest address: | Yes |
| IPv6 source address: | Yes |
| L4 Dest Port: | Yes |
| L4 Source Port: | Yes |

L2 settings:

| | |
|---------------------|-----|
| Dest Mac address: | No |
| Source Mac address: | No |
| Vlan Id: | Yes |
| Inner-vlan Id: | No |
| Incoming port: | Yes |

GRE settings:

| | |
|-----------|----|
| Key: | No |
| Protocol: | No |

MPLS settings:

| | |
|---------------|-----|
| MPLS Enabled: | Yes |
|---------------|-----|

VXLAN settings:

| | |
|-------------|----|
| VXLAN VNID: | No |
|-------------|----|

CHAPTER 25

Uplink Failure Detection Operational Commands

- `show uplink-failure-detection`

show uplink-failure-detection

| | |
|---------------------------------|---|
| Syntax | <code>show uplink-failure-detection</code> <code><group group-name></code> |
| Release Information | Command introduced in Junos OS Release 11.1 for EX Series switches. |
| Description | Display information about the uplink-failure-detection group, the member interfaces, and their status. |
| Options | none —Display information about all groups configured for uplink failure detection. group group-name —(Optional) Display information about the specified group only. |
| Required Privilege Level | view |
| Related Documentation | <ul style="list-style-type: none"> • Overview of Uplink Failure Detection on page 175 • Configuring Interfaces for Uplink Failure Detection on page 177 • Example: Configuring Interfaces for Uplink Failure Detection on page 178 |
| List of Sample Output | show uplink-failure-detection on page 446 |
| Output Fields | Table 47 on page 446 lists the output fields for the show uplink-failure-detection command. Output fields are listed in the approximate order in which they appear. |

Table 47: show uplink-failure-detection Output Fields

| Field Name | Field Description |
|----------------|---|
| Group | Name of the group. |
| Uplink | The uplink interface or interfaces configured as link-to-monitor. NOTE: The asterisk (*) indicates that the link is up. |
| Downlink | The downlink interface or interfaces configured as link-to-disable. NOTE: The asterisk (*) indicates that the link is up. |
| Failure Action | Status of uplink failure detection: <ul style="list-style-type: none"> • Active—The switch has detected an uplink failure and has brought the downlink down. • Inactive—The uplink or uplinks are up. |

Sample Output

show uplink-failure-detection

```
user@switch> show uplink-failure-detection
```

Group : group1
Uplink : ge-0/0/0*
Downlink : ge-0/0/1*
Failure Action : Inactive

Group : group2
Uplink : ge-0/0/3.0
Downlink : ge-0/0/4.0
Failure Action : Active

